



US009405706B2

(12) **United States Patent**  
**Doshi et al.**

(10) **Patent No.:** **US 9,405,706 B2**  
(45) **Date of Patent:** **Aug. 2, 2016**

(54) **INSTRUCTION AND LOGIC FOR ADAPTIVE DATASET PRIORITIES IN PROCESSOR CACHES**

(71) Applicant: **Intel Corporation**, Santa Clara, CA (US)

(72) Inventors: **Kshitij A. Doshi**, Chandler, AZ (US);  
**Karthik Raman**, Hillsboro, OR (US);  
**Christopher J. Hughes**, Santa Clara, CA (US)

(73) Assignee: **Intel Corporation**, Santa Clara, CA (US)

(\*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 115 days.

(21) Appl. No.: **14/496,255**

(22) Filed: **Sep. 25, 2014**

(65) **Prior Publication Data**

US 2016/0092373 A1 Mar. 31, 2016

(51) **Int. Cl.**  
**G06F 12/12** (2016.01)  
**G06F 12/08** (2016.01)

(52) **U.S. Cl.**  
CPC ..... **G06F 12/126** (2013.01); **G06F 12/0875** (2013.01); **G06F 12/0897** (2013.01); **G06F 2212/452** (2013.01); **G06F 2212/69** (2013.01)

(58) **Field of Classification Search**  
CPC ..... G06F 12/12; G06F 12/08  
See application file for complete search history.

(56) **References Cited**

U.S. PATENT DOCUMENTS

5,774,685 A \* 6/1998 Dubey ..... G06F 9/383  
711/122  
6,532,520 B1 \* 3/2003 Dean ..... G06F 12/121  
711/123  
7,899,994 B2 \* 3/2011 Zhao ..... G06F 13/00  
711/133

2006/0161762 A1 \* 7/2006 Eisen ..... G06F 9/3844  
712/233  
2006/0190924 A1 \* 8/2006 Bruening ..... G06F 12/0875  
717/104  
2007/0079073 A1 \* 4/2007 Rosenbluth ..... G06F 9/30047  
711/134  
2007/0123470 A1 \* 5/2007 Hogenhuis ..... A61K 31/426  
514/5.1  
2008/0025210 A1 \* 1/2008 Honary ..... H04L 47/10  
370/229  
2008/0120469 A1 \* 5/2008 Kornegay ..... G06F 12/122  
711/133  
2009/0210382 A1 \* 8/2009 Louzoun ..... G11C 15/00  
707/3  
2010/0064109 A1 \* 3/2010 Bull ..... G06F 12/126  
711/151  
2012/0227045 A1 \* 9/2012 Knauth ..... G06F 9/3863  
718/100  
2014/0089559 A1 \* 3/2014 Cai ..... G06F 12/126  
711/102  
2015/0067691 A1 \* 3/2015 Johnson ..... G06F 9/5016  
718/103  
2015/0186280 A1 \* 7/2015 Lepak ..... G06F 12/121  
711/136

#### OTHER PUBLICATIONS

Park et al., "Location-Aware Cache Management for Many-Core Processors with Deep Cache Hierarchy", SC '13, 12 pages, Nov. 2013.

Schmidt and Cranor, "Half-Sync/Half-Async: An Architectural Pattern for Efficient and Well-structured Concurrent I/O", Proc. Pattern Languages of Programs, 11 pages, 1995.

\* cited by examiner

Primary Examiner — Adam M Queler

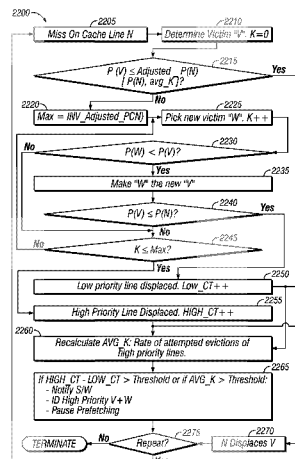
Assistant Examiner — Kenneth Tsang

(74) Attorney, Agent, or Firm — Baker Botts L.L.P.

(57) **ABSTRACT**

A processor includes a front end, a cache, and a cache controller. The front end includes logic to receive an instruction defining a priority dataset. The priority dataset includes ranges of memory addresses each corresponding to a respective priority level. The cache controller includes logic to detect a miss in the cache for a requested cache value, determine a candidate cache victim from the cache, determine a priority of the requested cache value and the candidate cache victim according to the priority dataset, and evict the candidate cache victim based on a determination that the priority of the candidate cache victim is less or equal to the priority of the requested cache value.

**12 Claims, 30 Drawing Sheets**



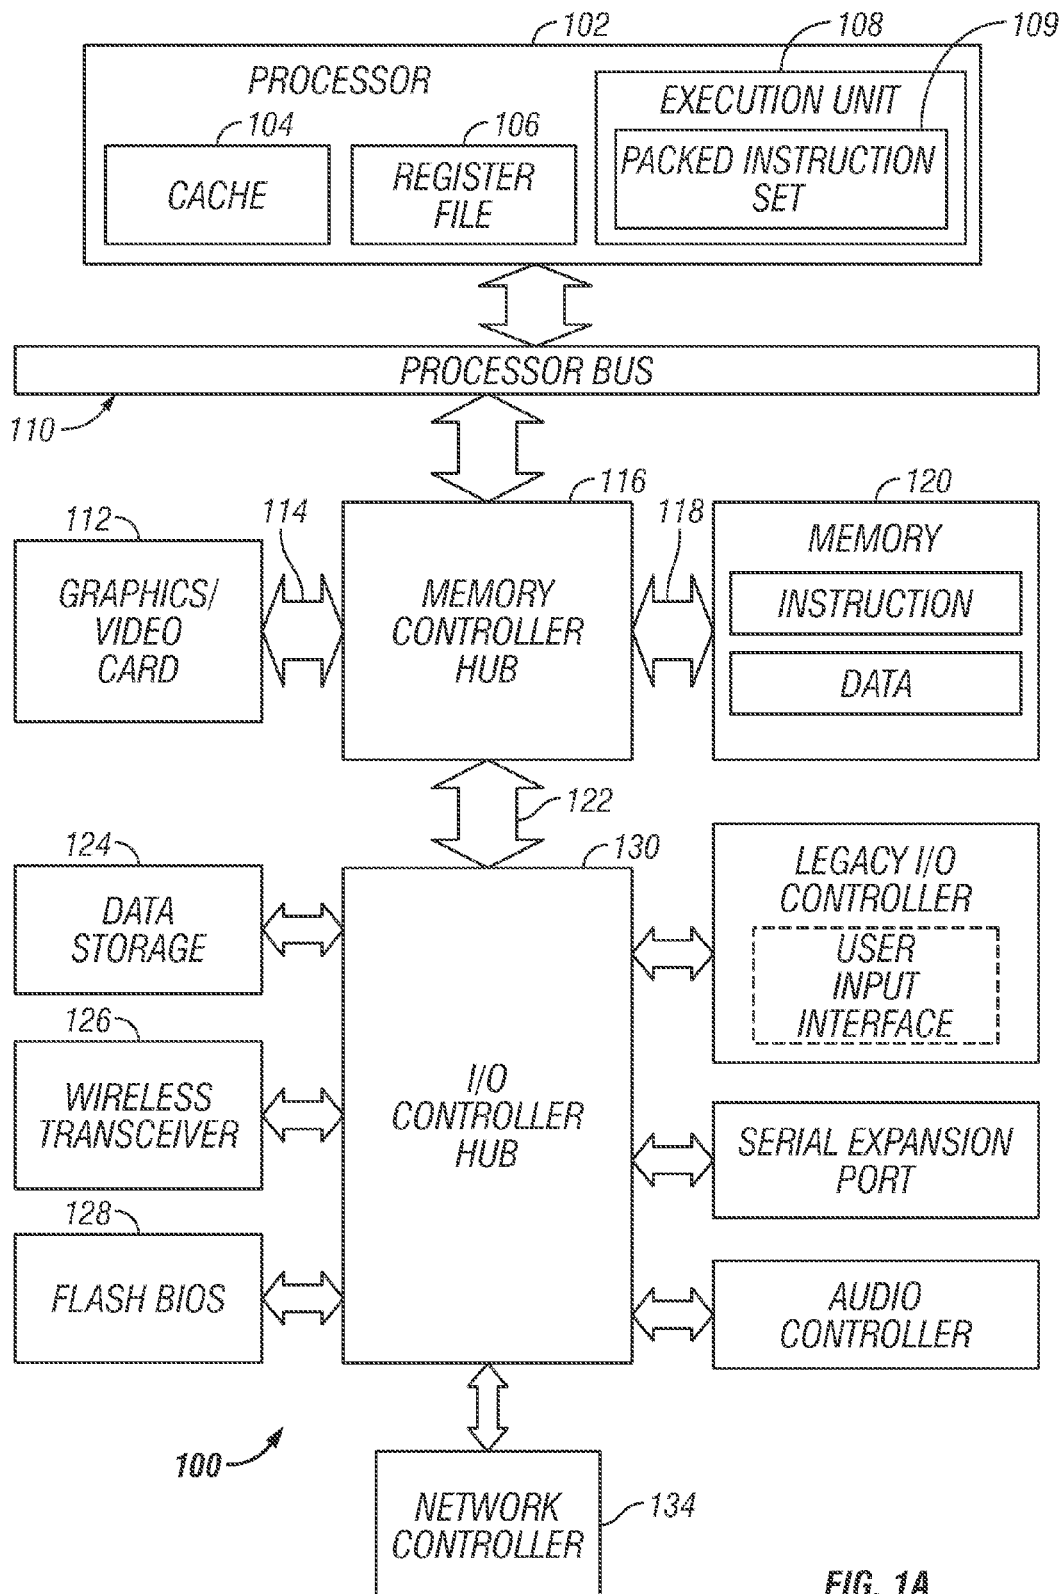


FIG. 1A

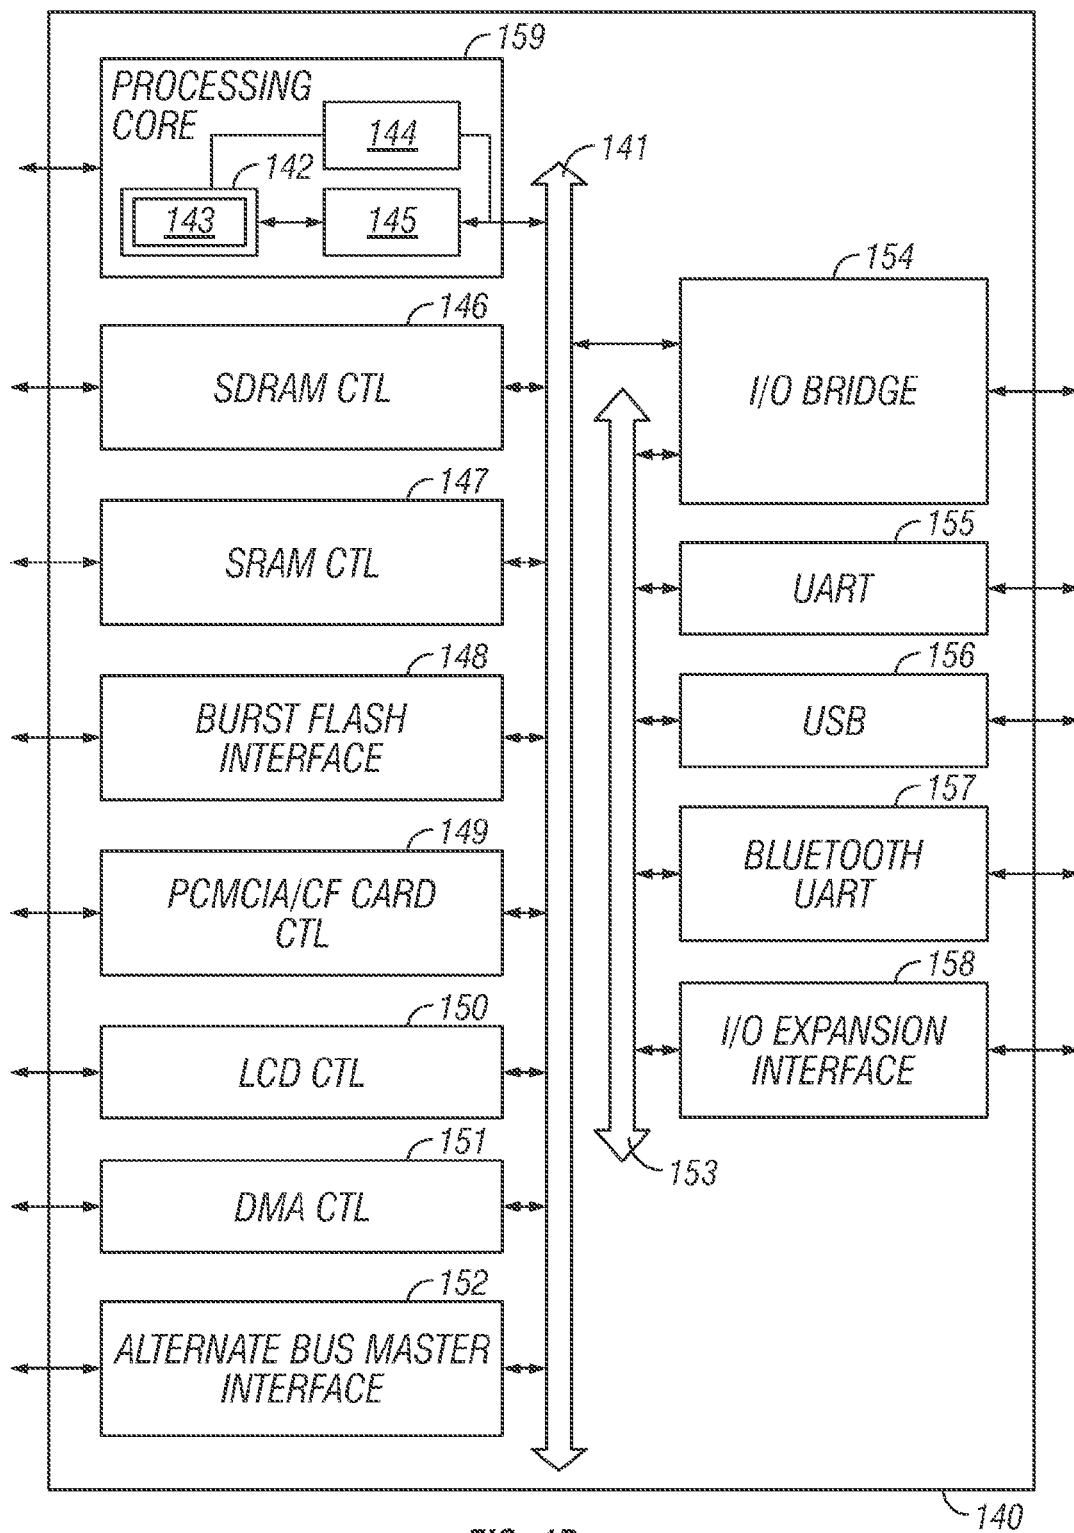


FIG. 1B

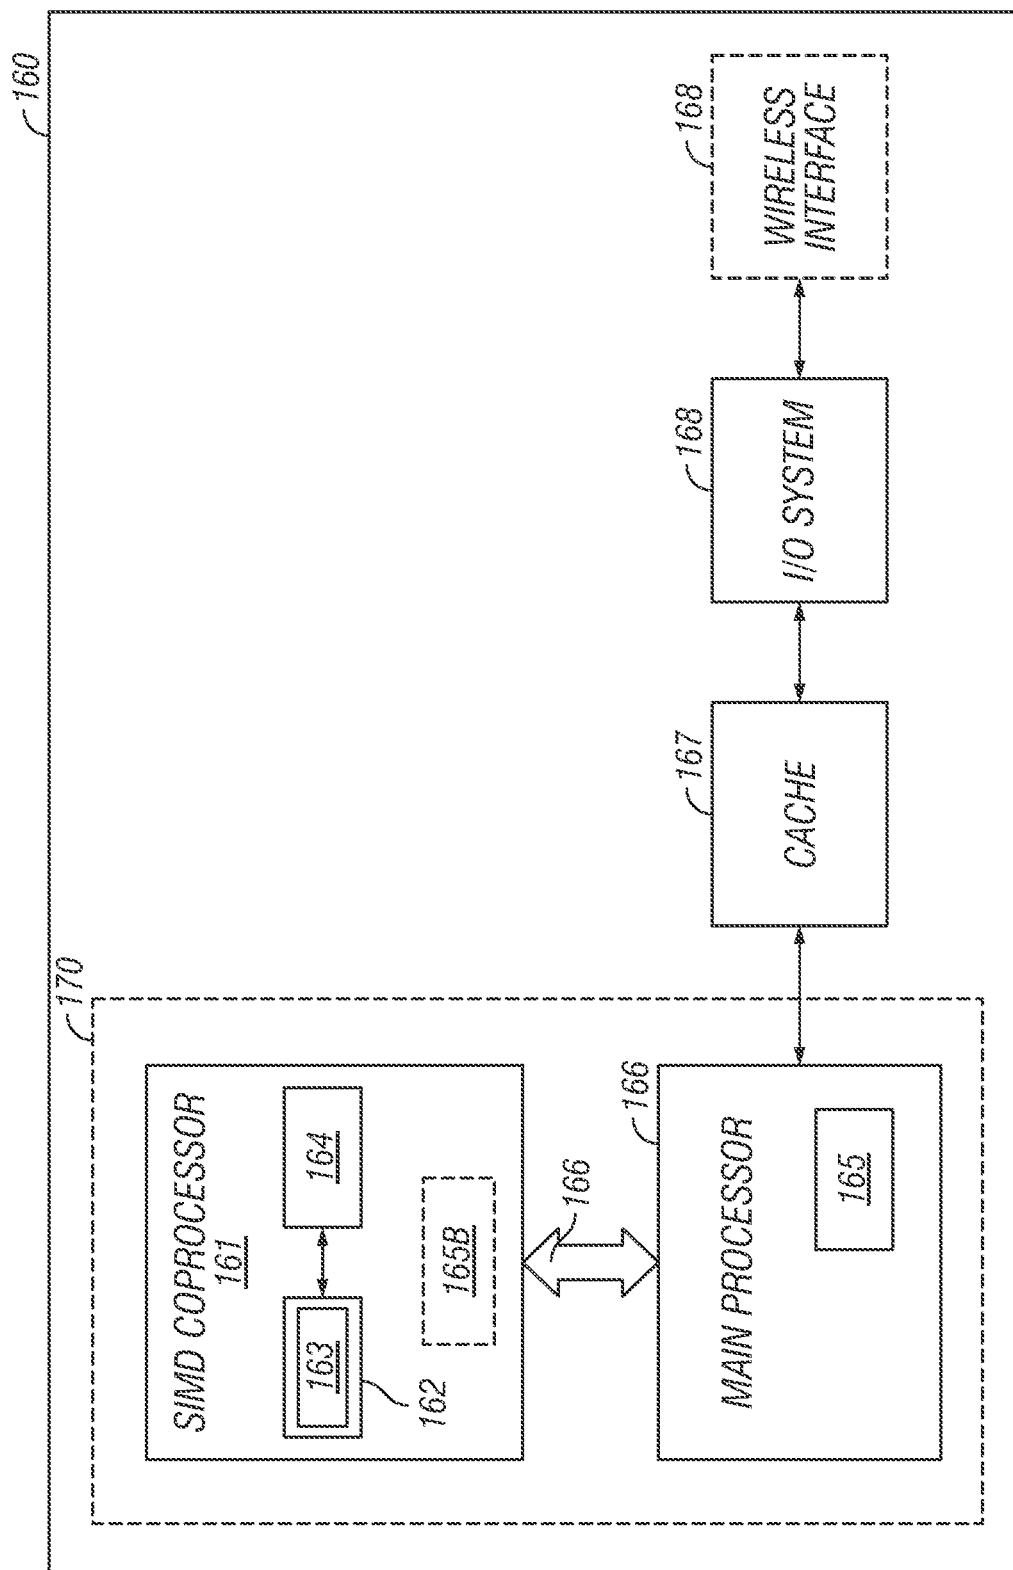


FIG. 1C

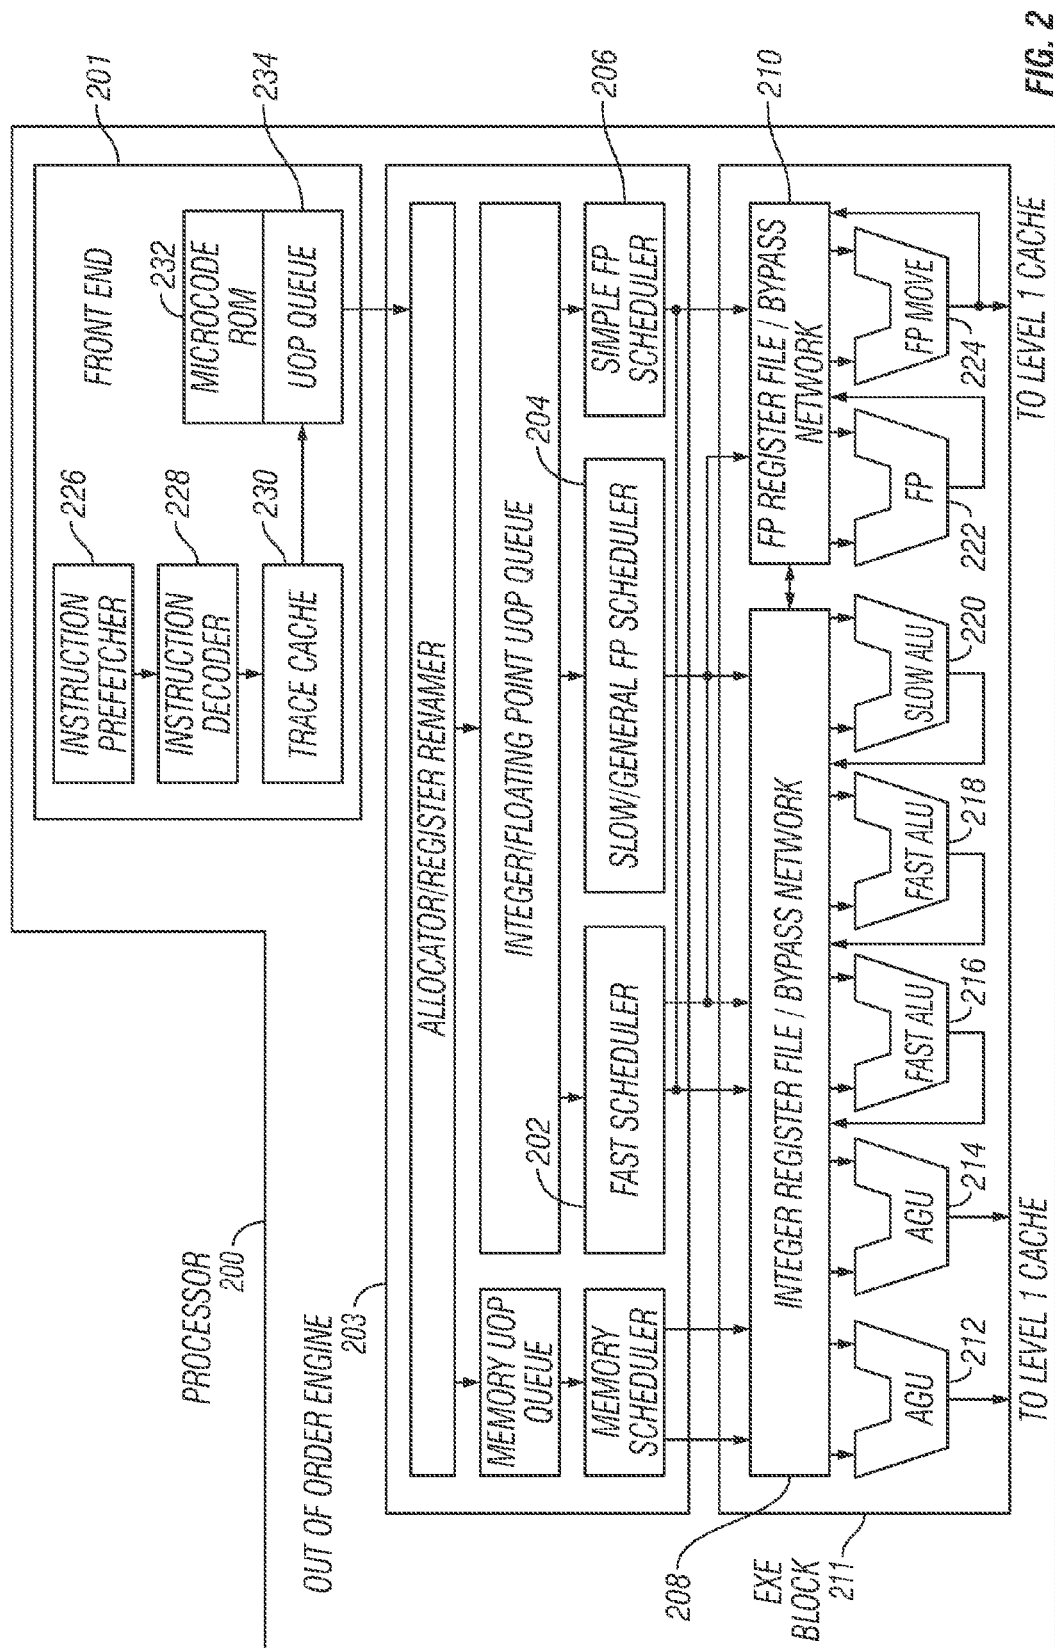


FIG. 2

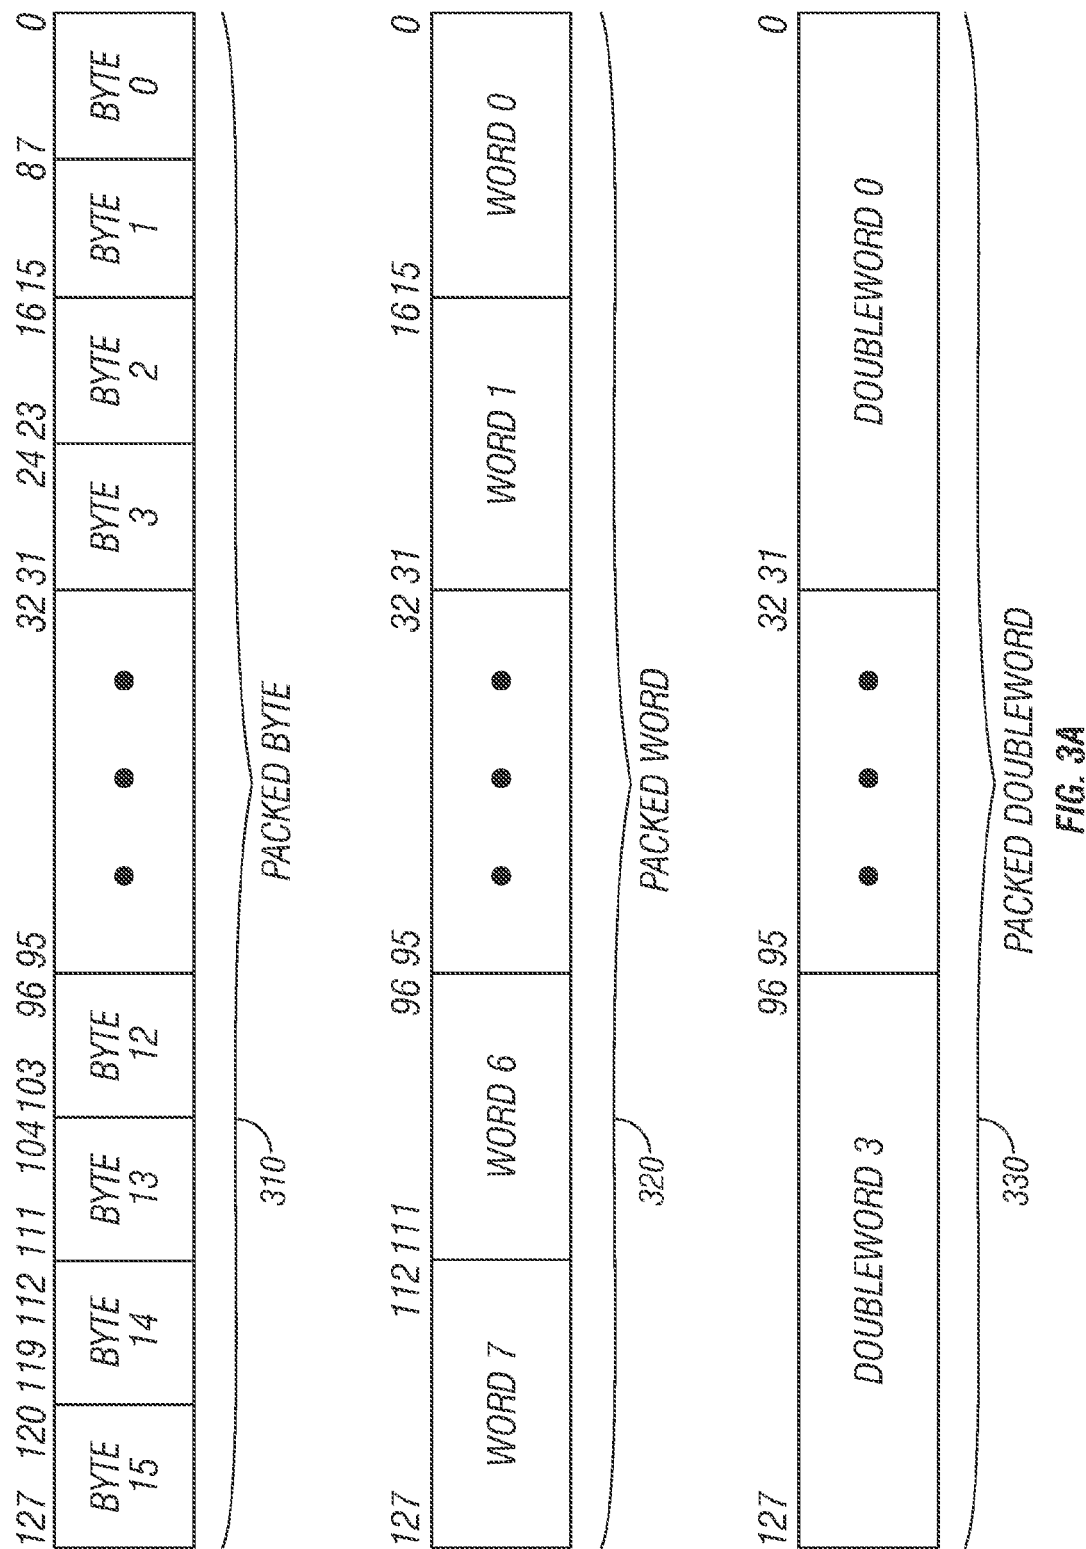


FIG. 3A

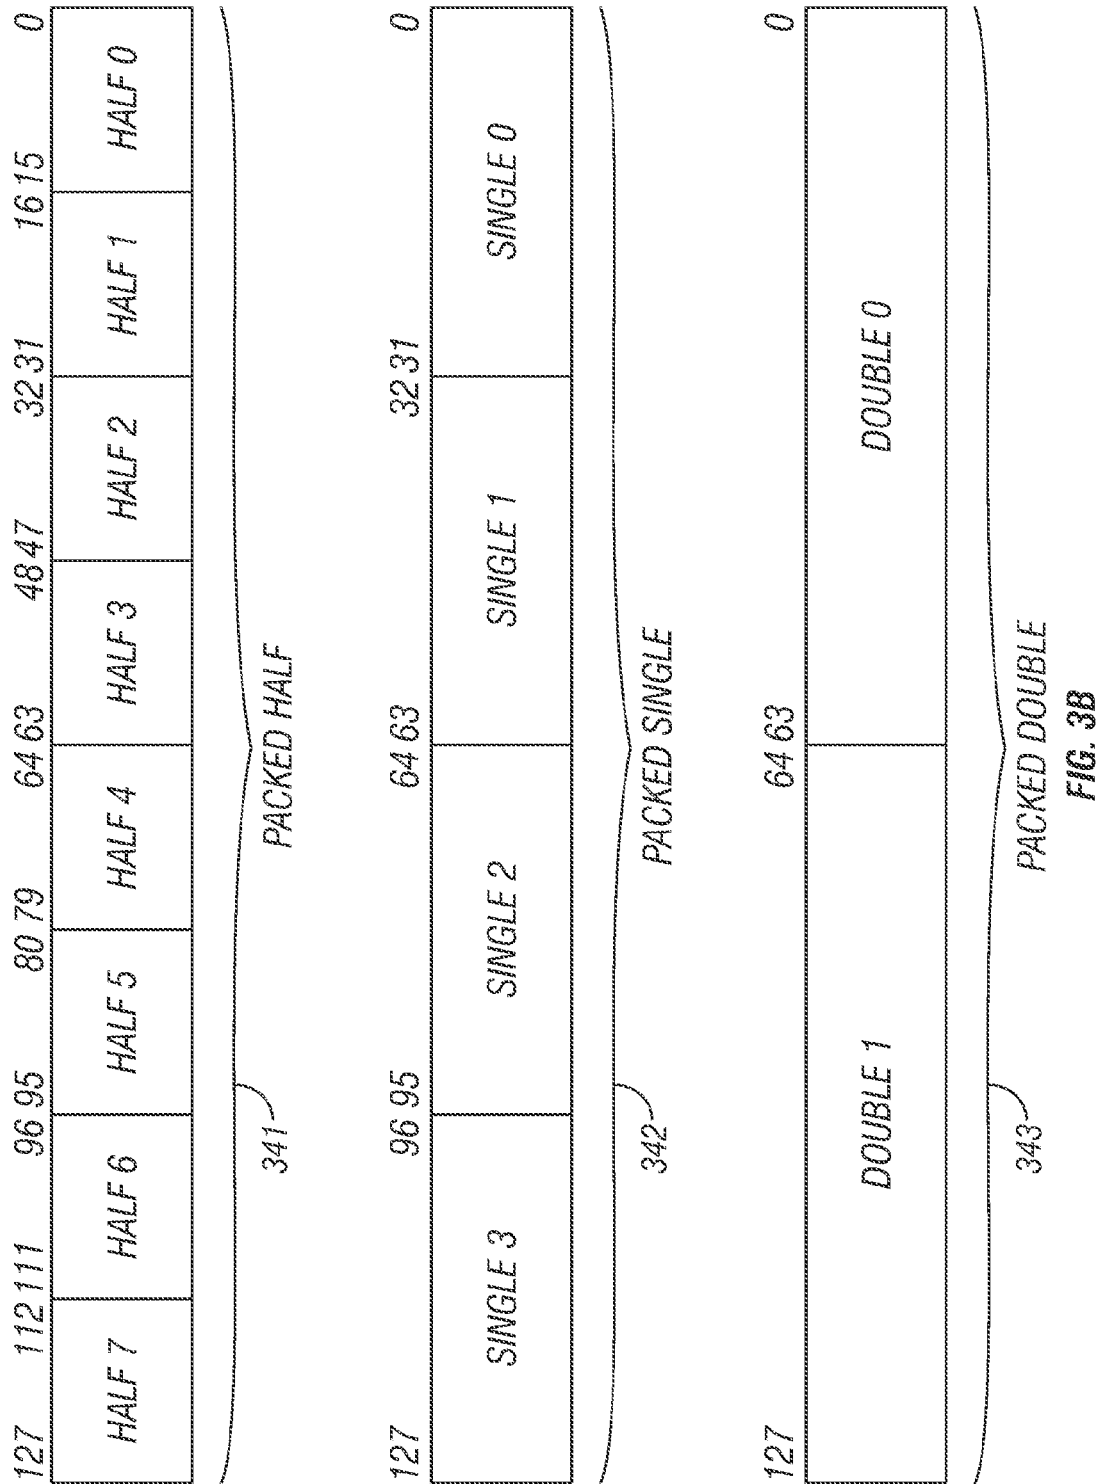
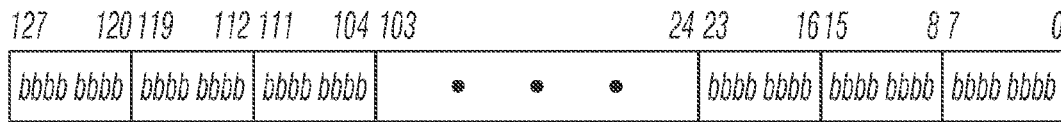
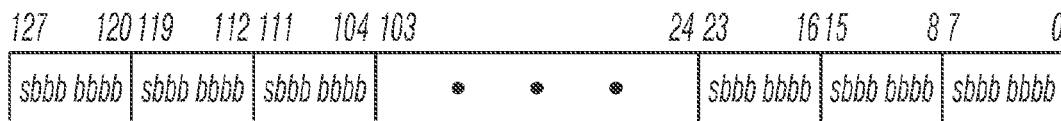


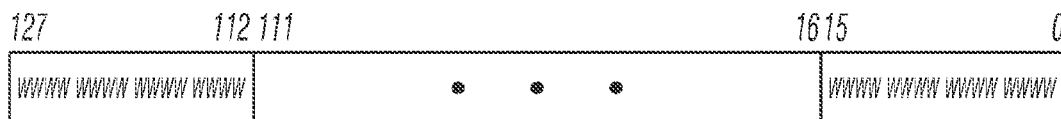
FIG. 3B



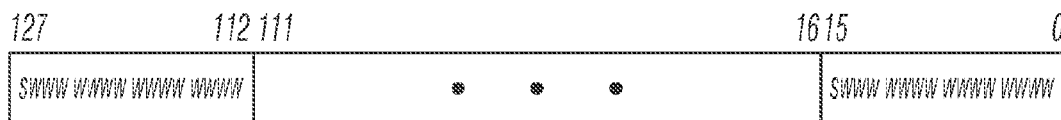
UNSIGNED PACKED BYTE REPRESENTATION 344



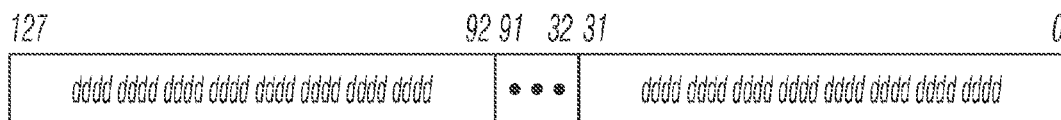
SIGNED PACKED BYTE REPRESENTATION 345



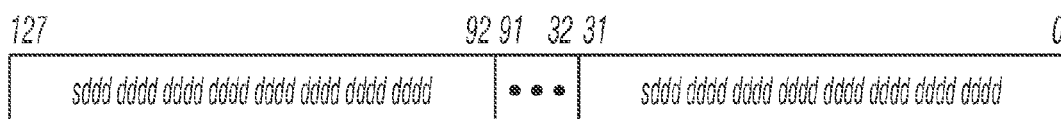
UNSIGNED PACKED WORD REPRESENTATION 346



SIGNED PACKED WORD REPRESENTATION 347



UNSIGNED PACKED DOUBLEWORD REPRESENTATION 348



SIGNED PACKED DOUBLEWORD REPRESENTATION 349

FIG. 3C



31	1615	87	65	0
<u>361</u>	<u>362</u>	<u>363</u>	<u>364</u>	<u>365</u>

360

FIG. 3D

39	3231	1615	87	65	0
<u>378</u>	<u>371</u>	<u>372</u>	<u>373</u>	<u>374</u>	<u>375</u>

370

FIG. 3E

31	2827	2423	2221	2019	1615	1211	87	5	4	3	0
<u>381</u>	<u>382</u>	<u>383</u>	<u>384</u>	<u>385</u>	<u>386</u>	<u>387</u>	<u>388</u>	<u>389</u>	<u>390</u>		

380

FIG. 3F

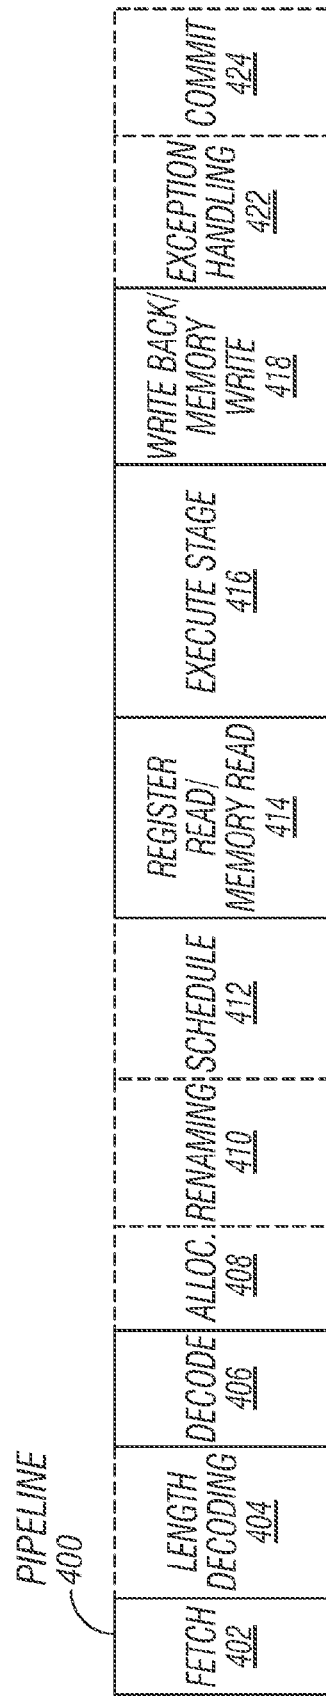


FIG. 4A

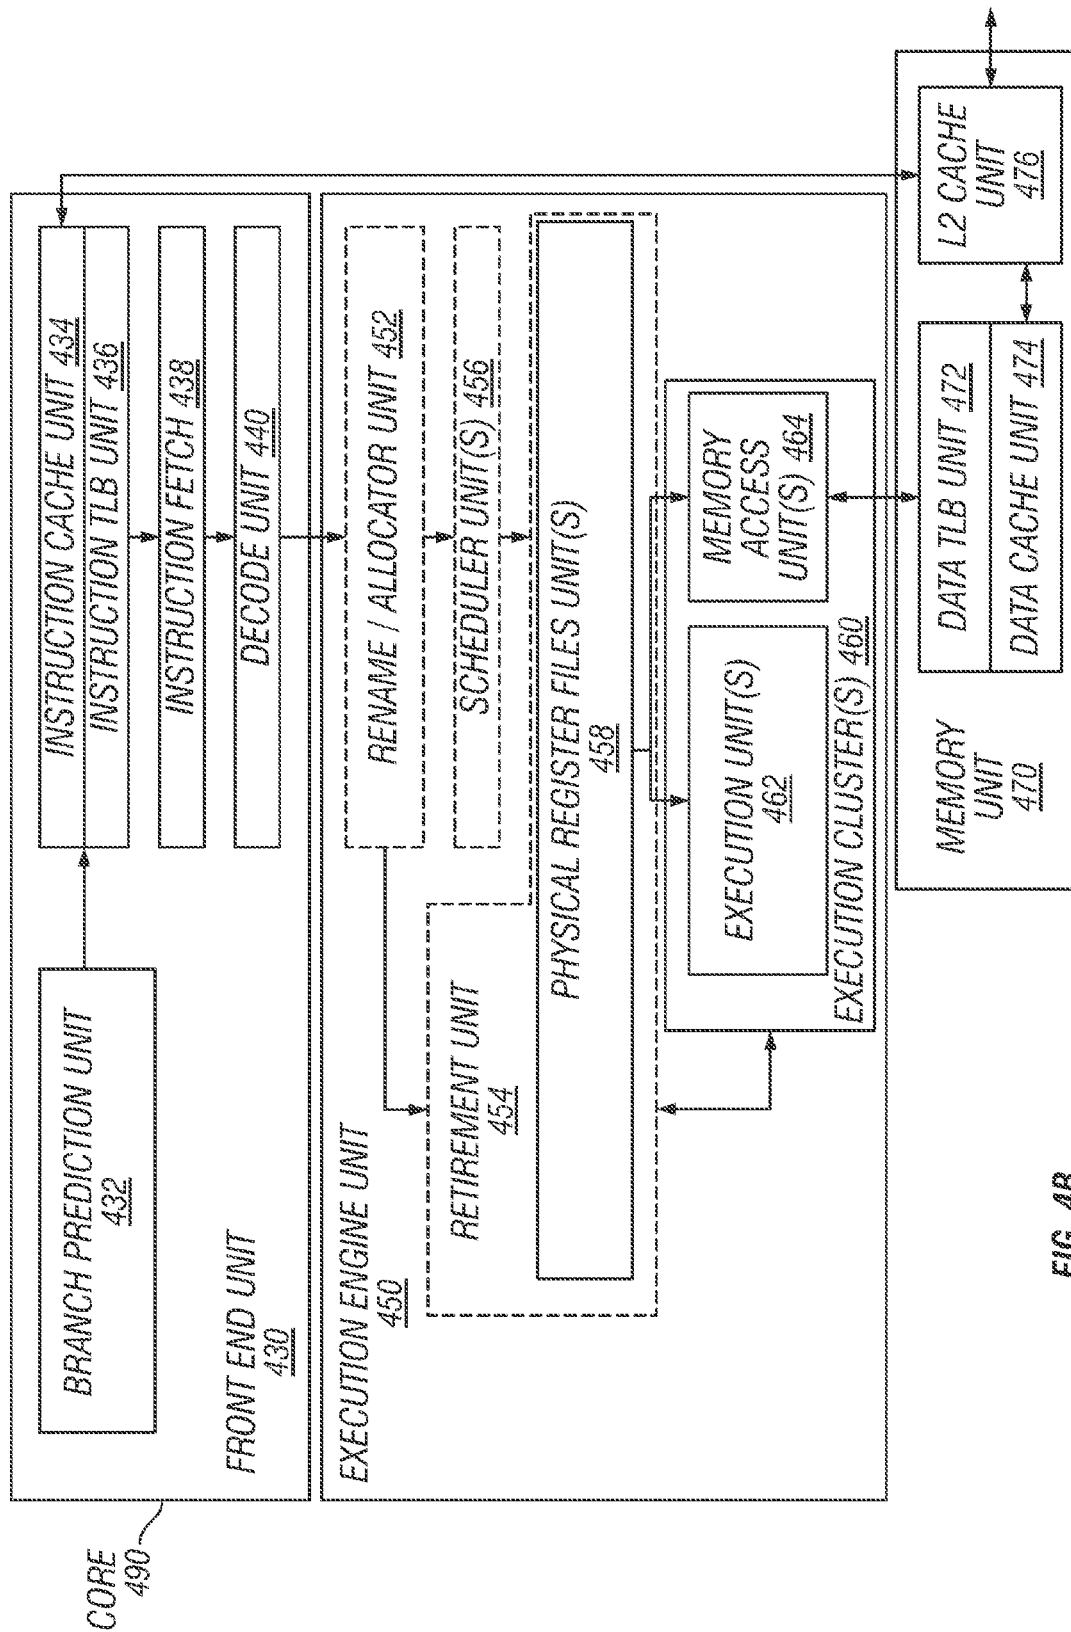


FIG. 4B

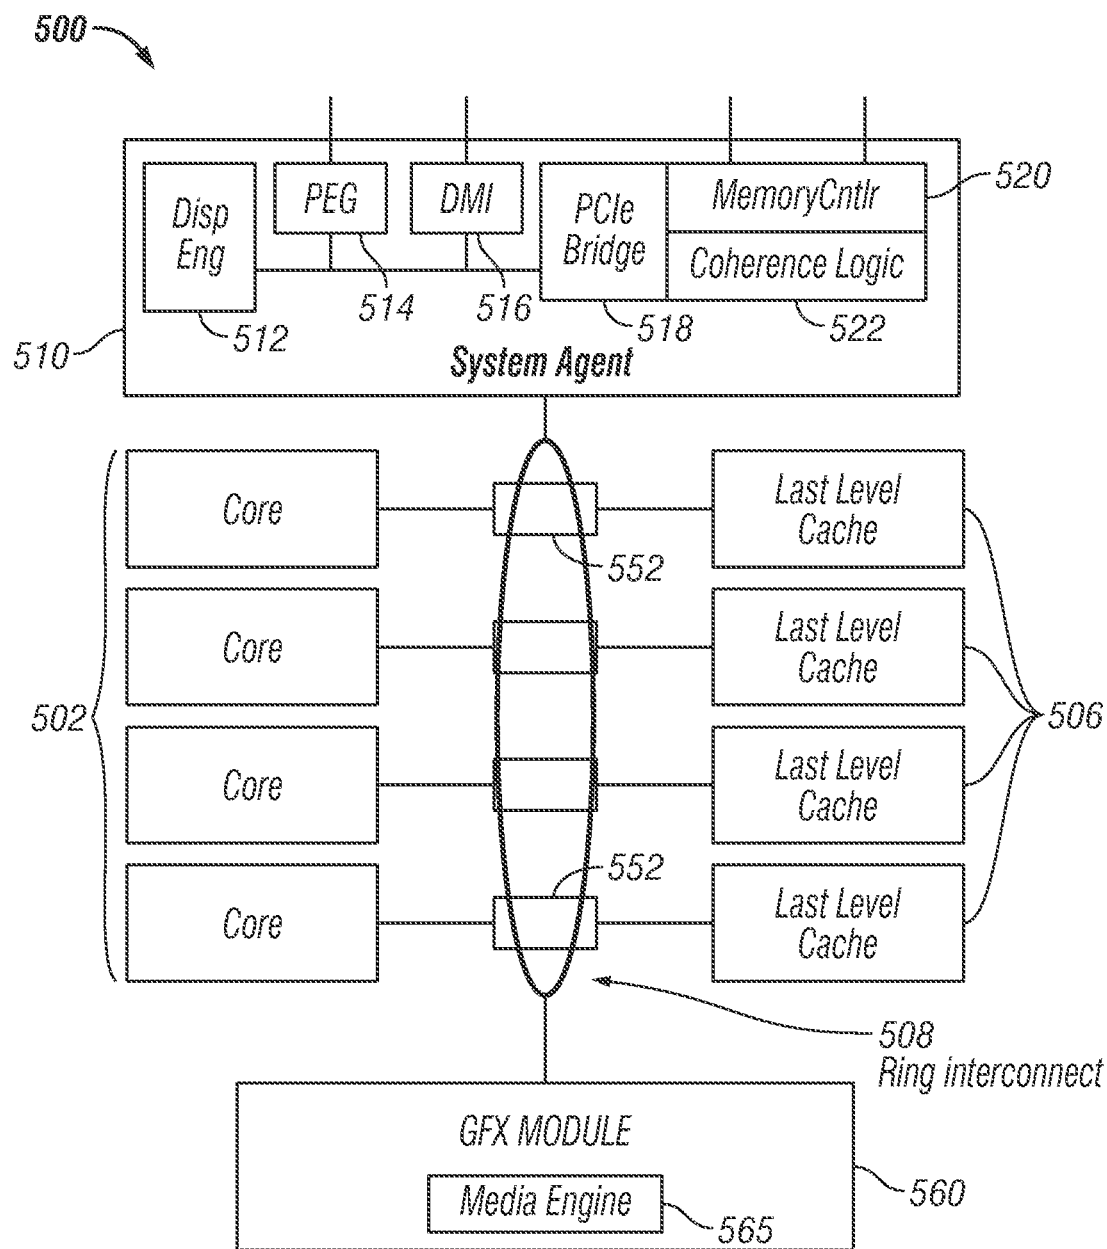


FIG. 5A

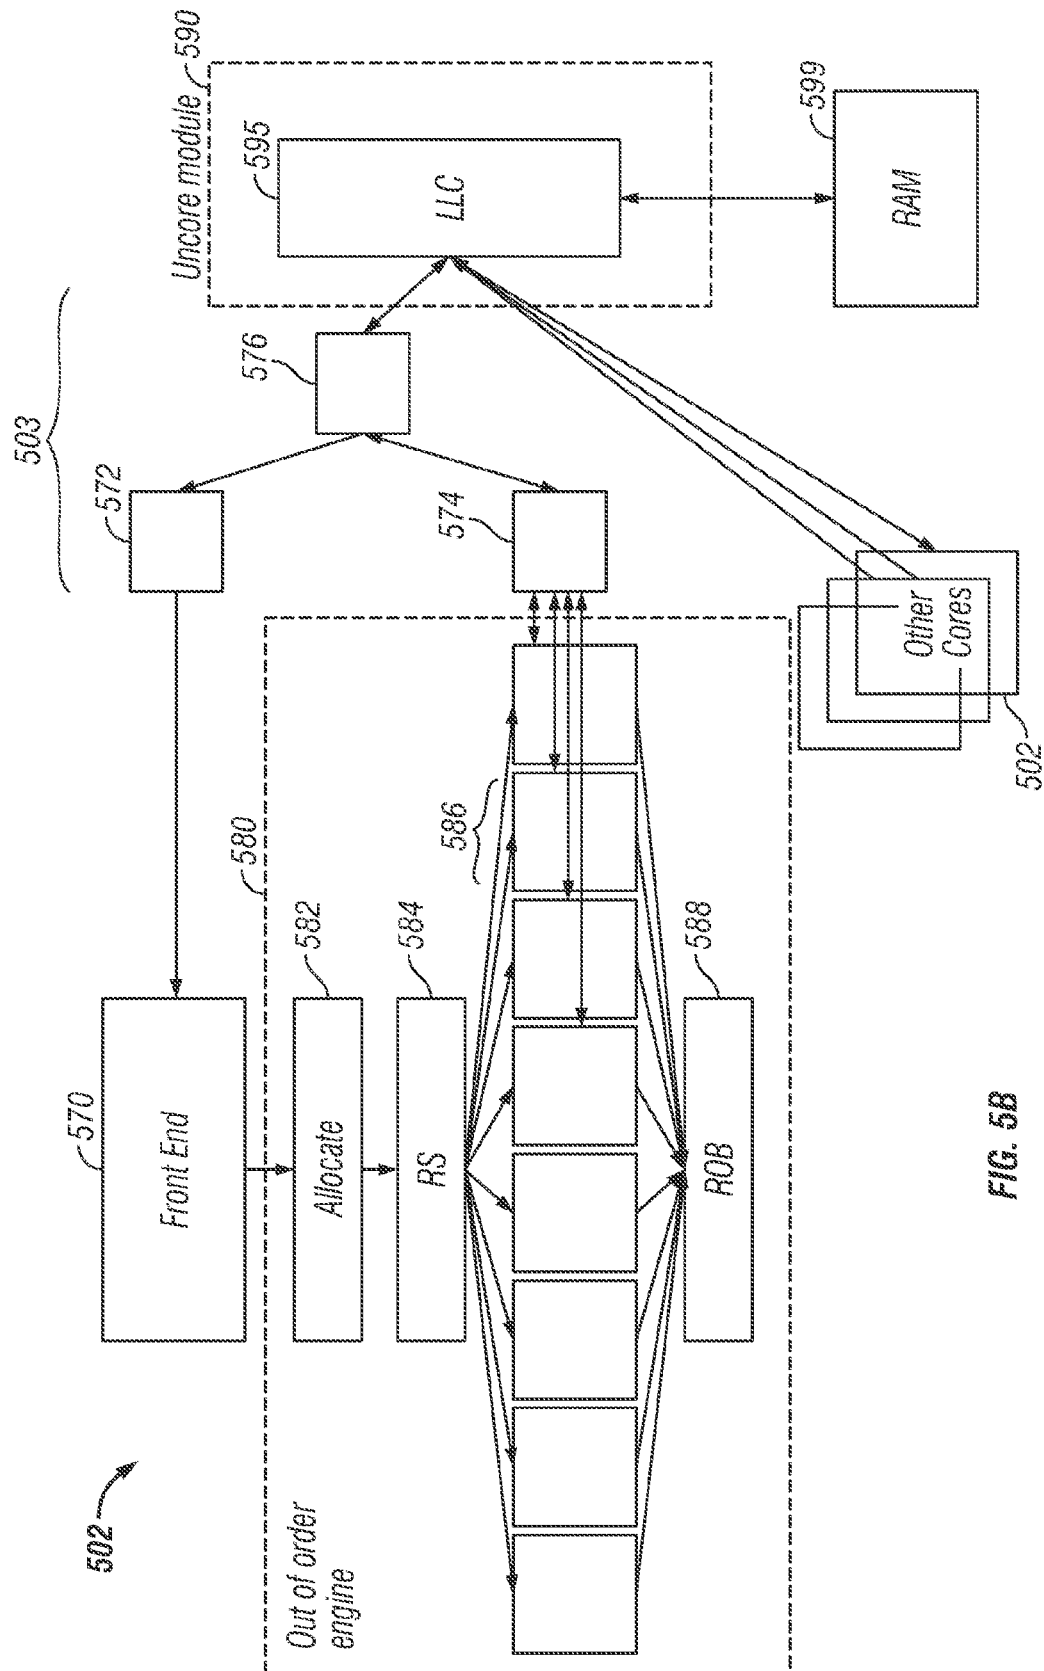


FIG. 5B

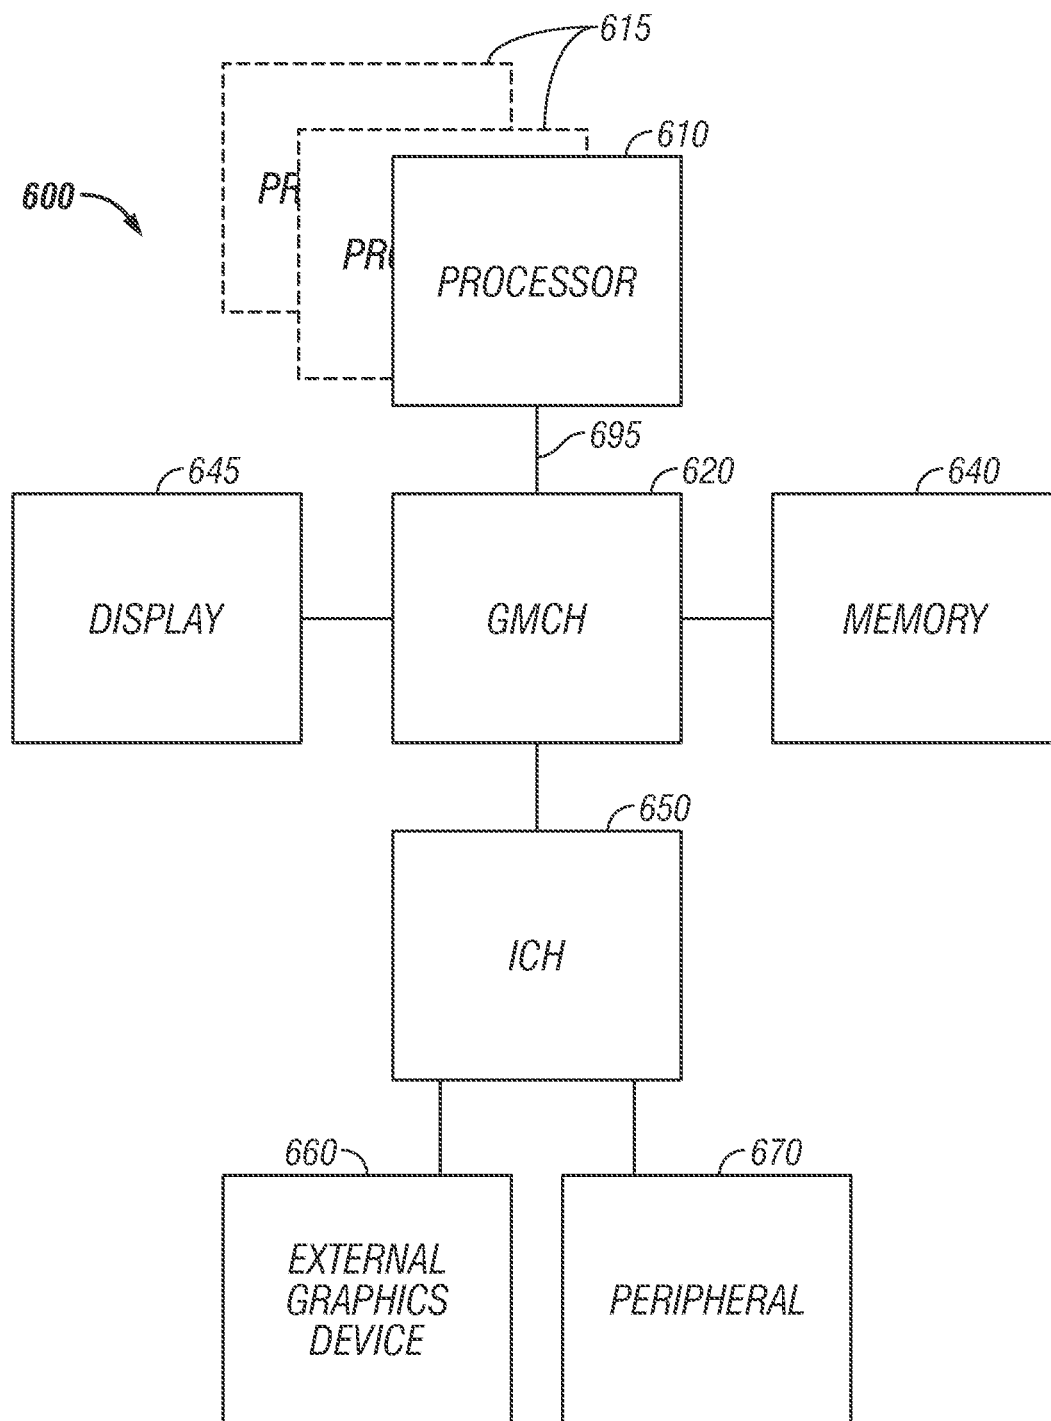


FIG. 6

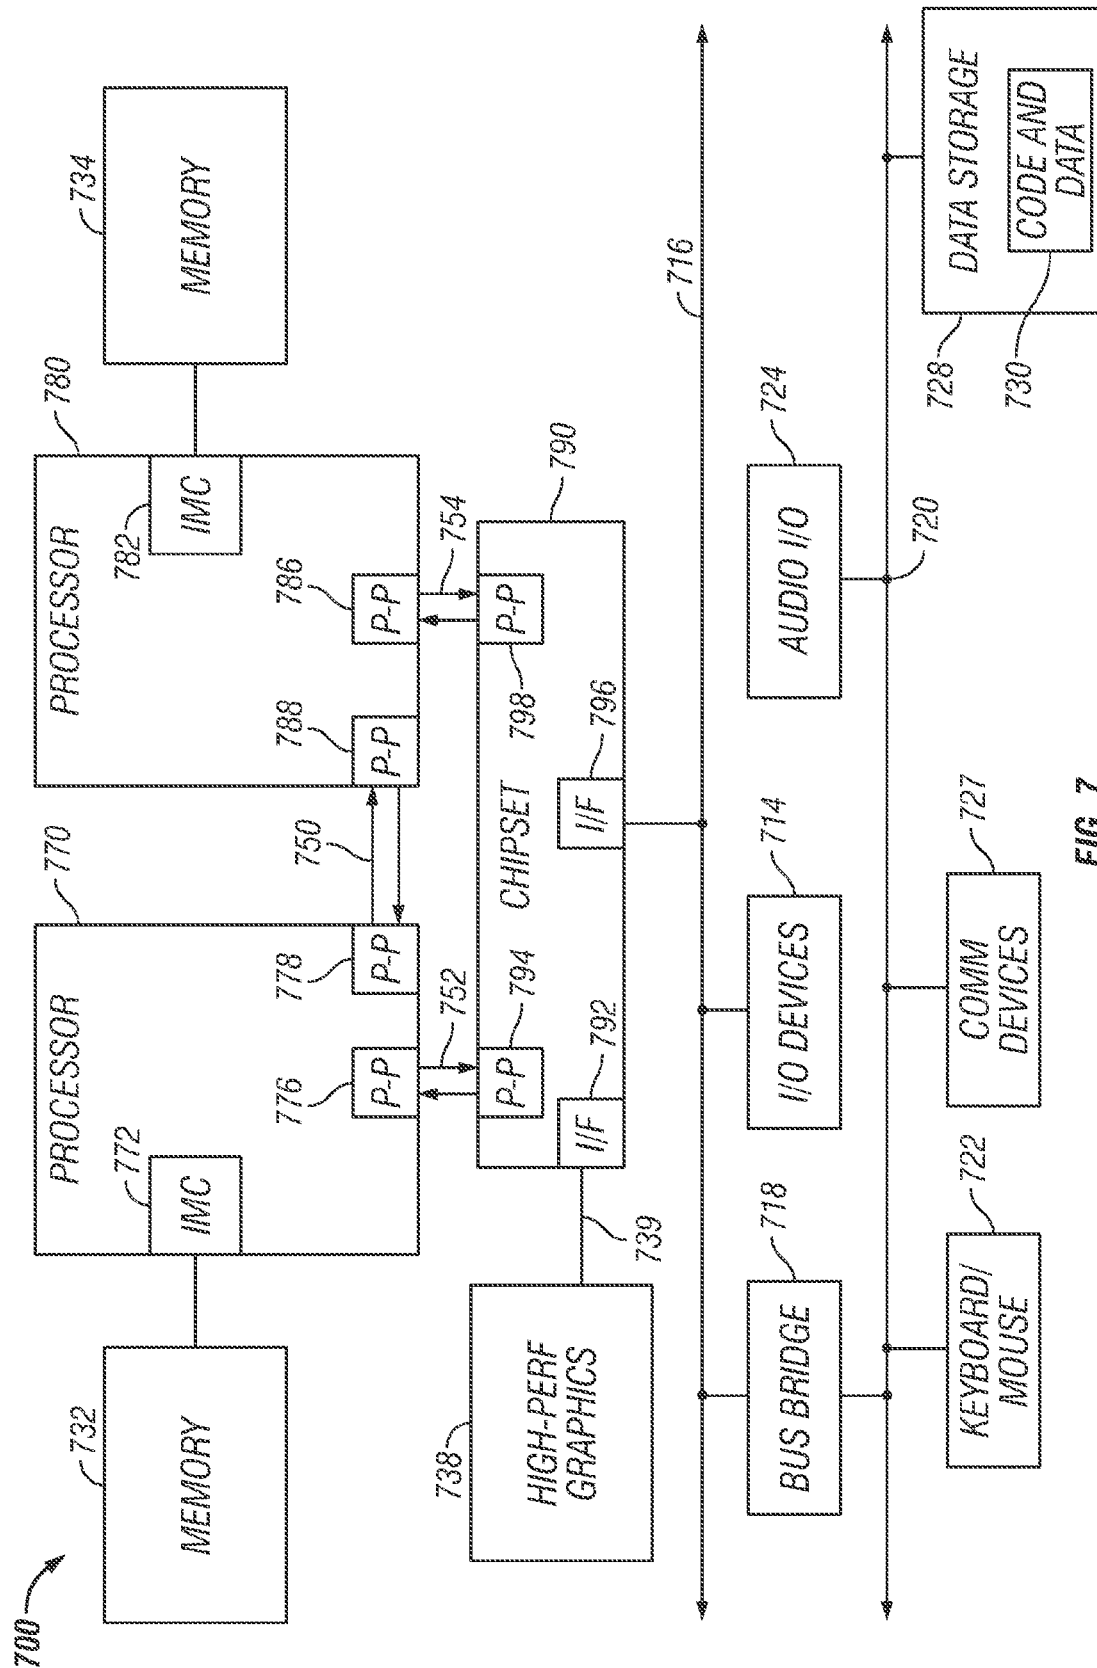


FIG. 7

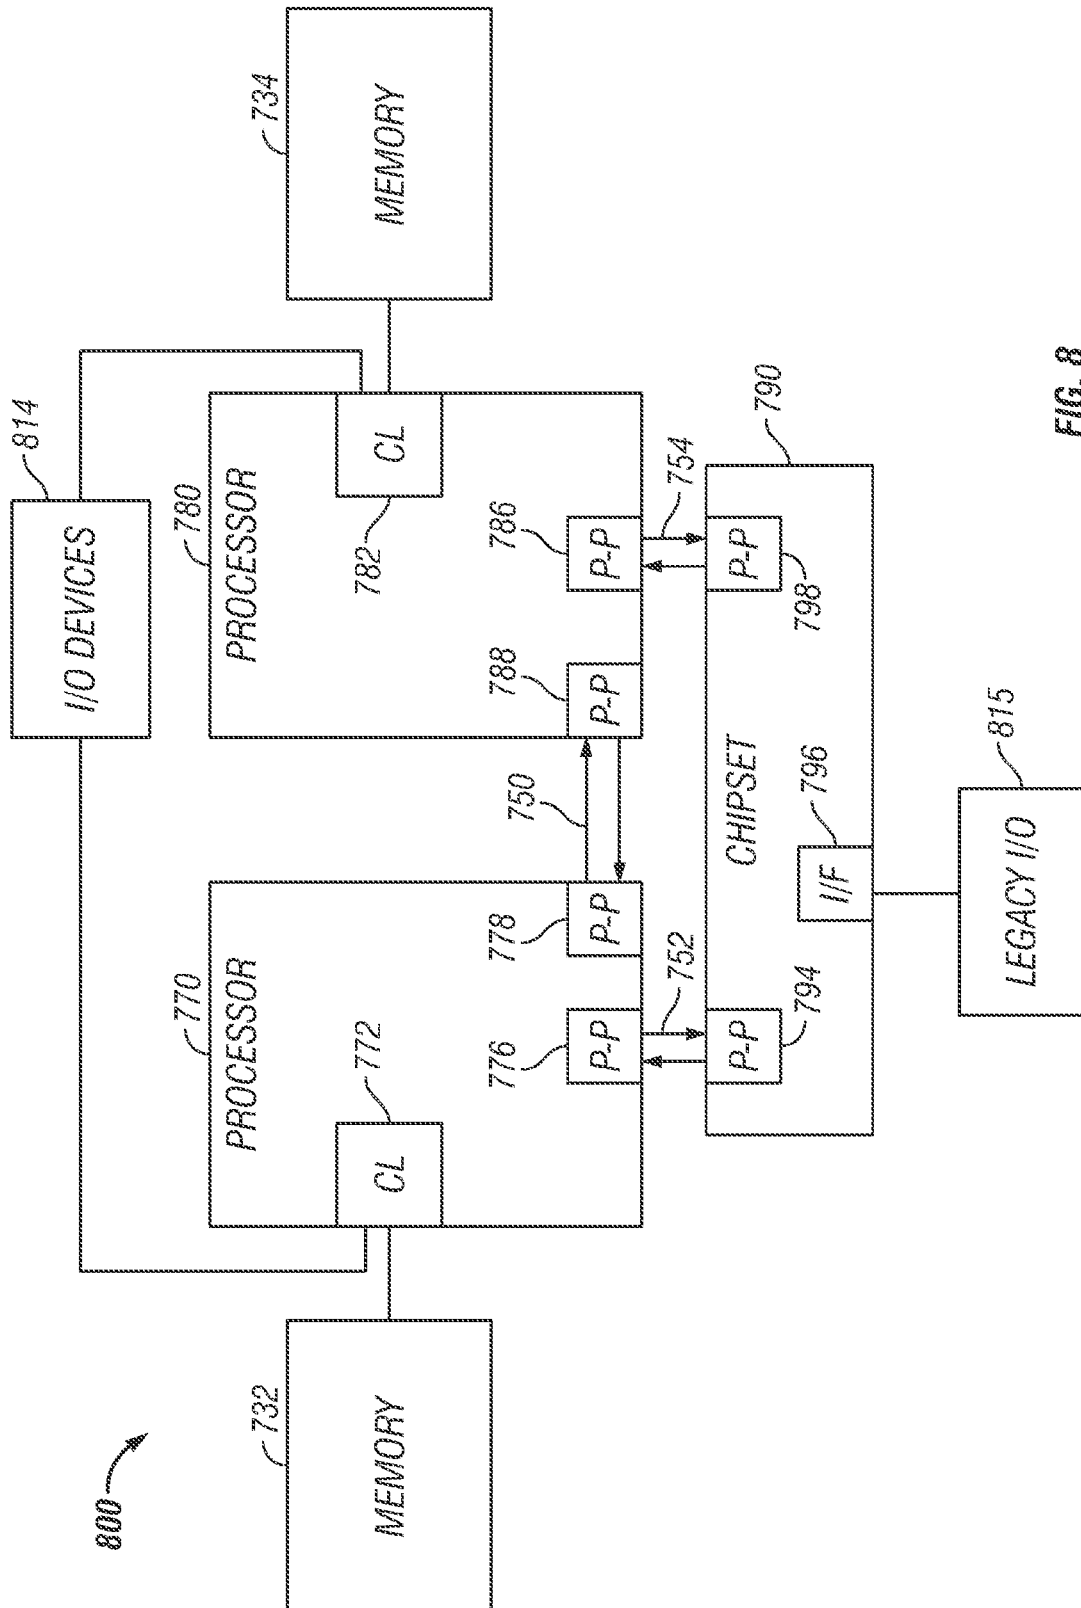


FIG. 8



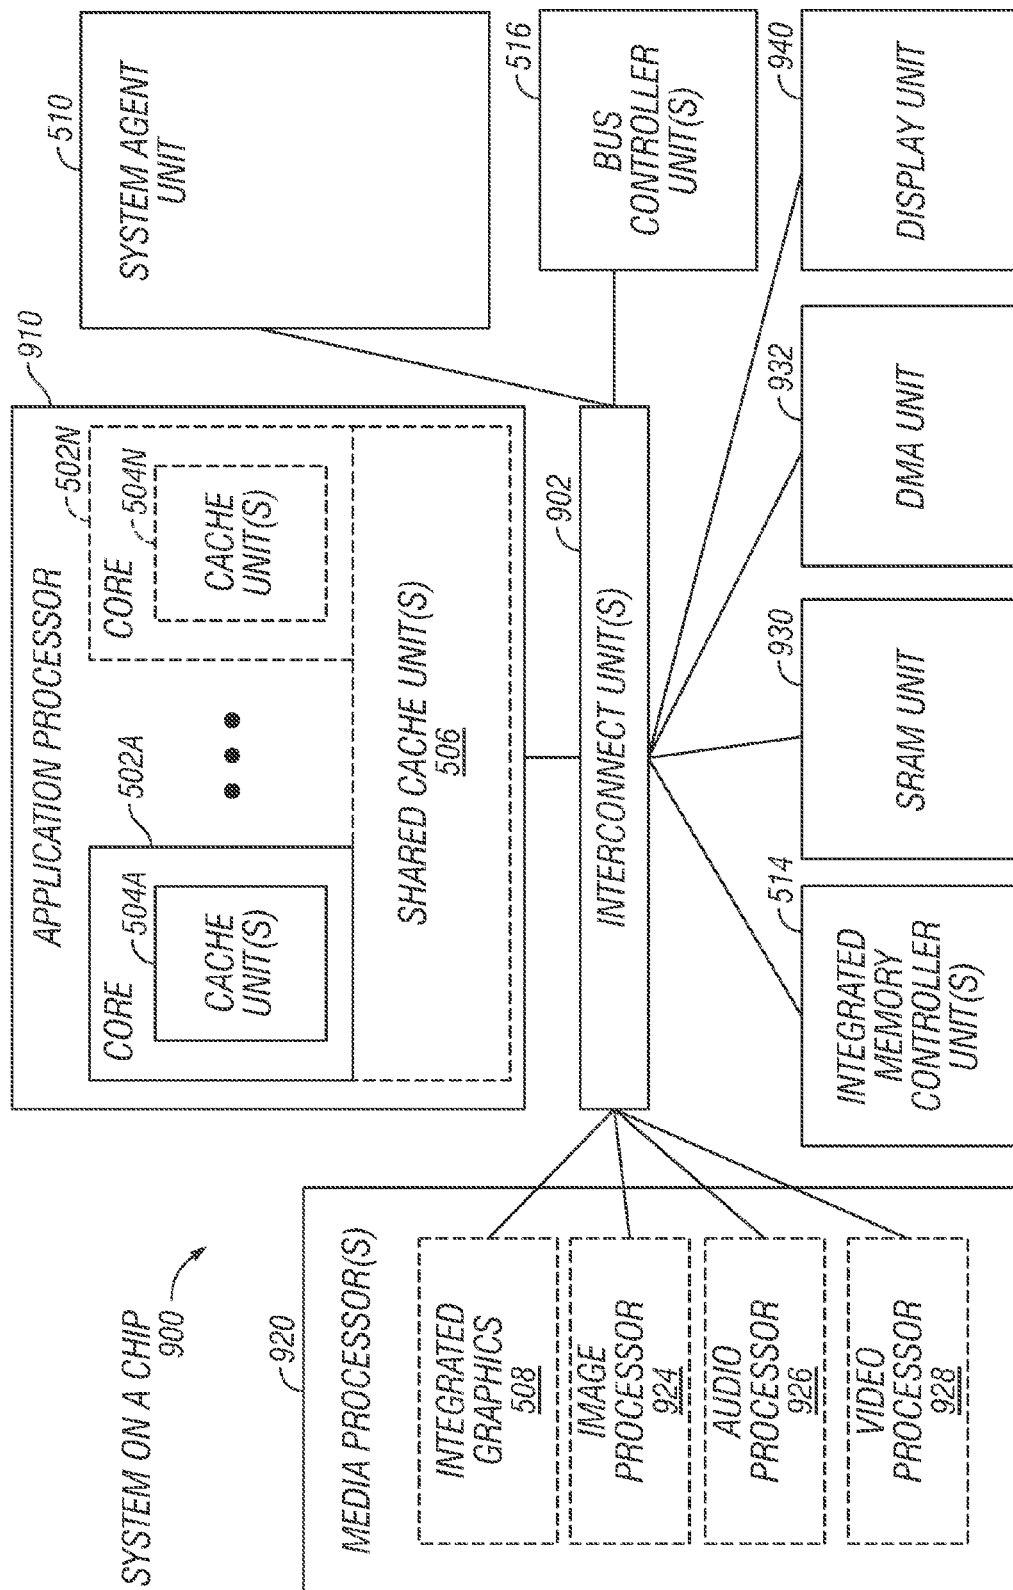


FIG. 9

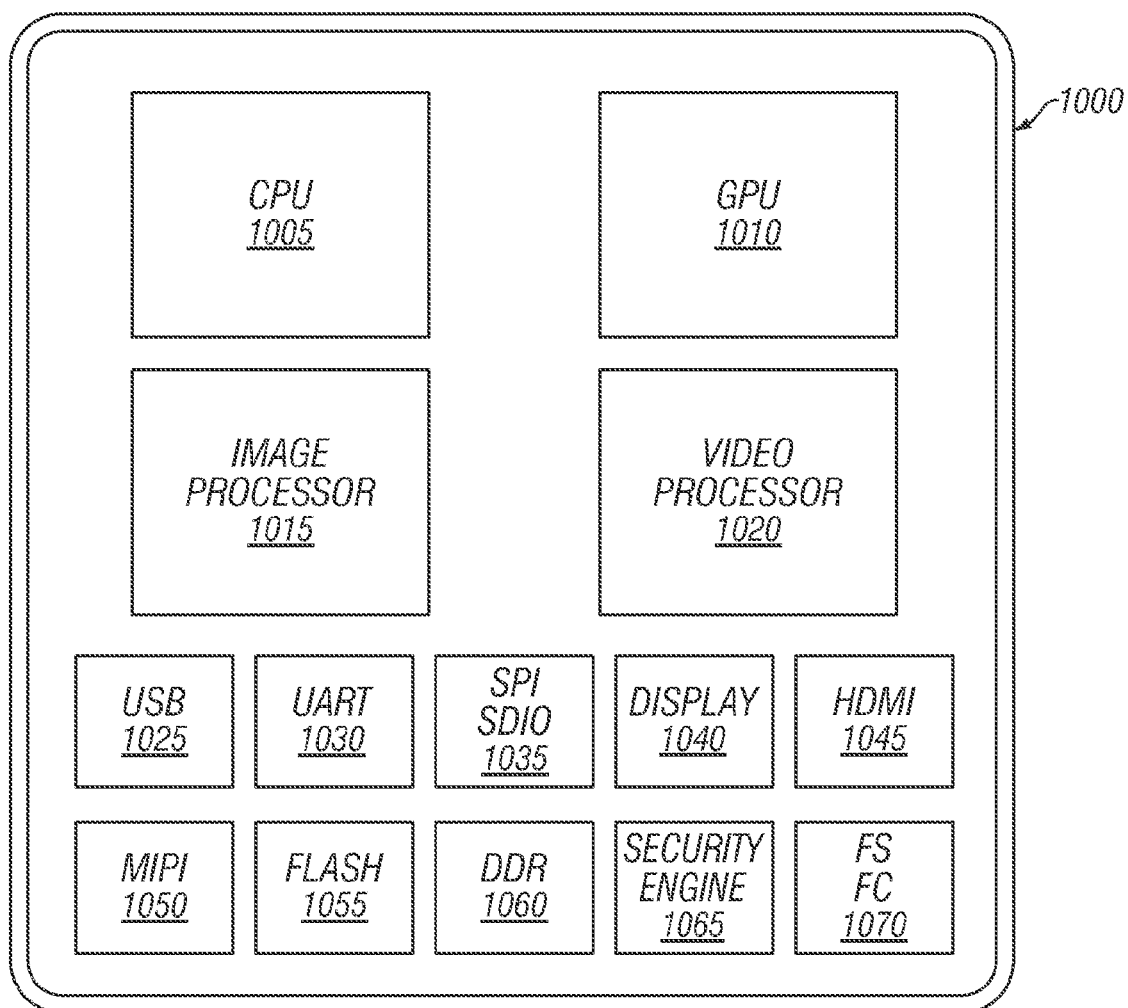


FIG. 10

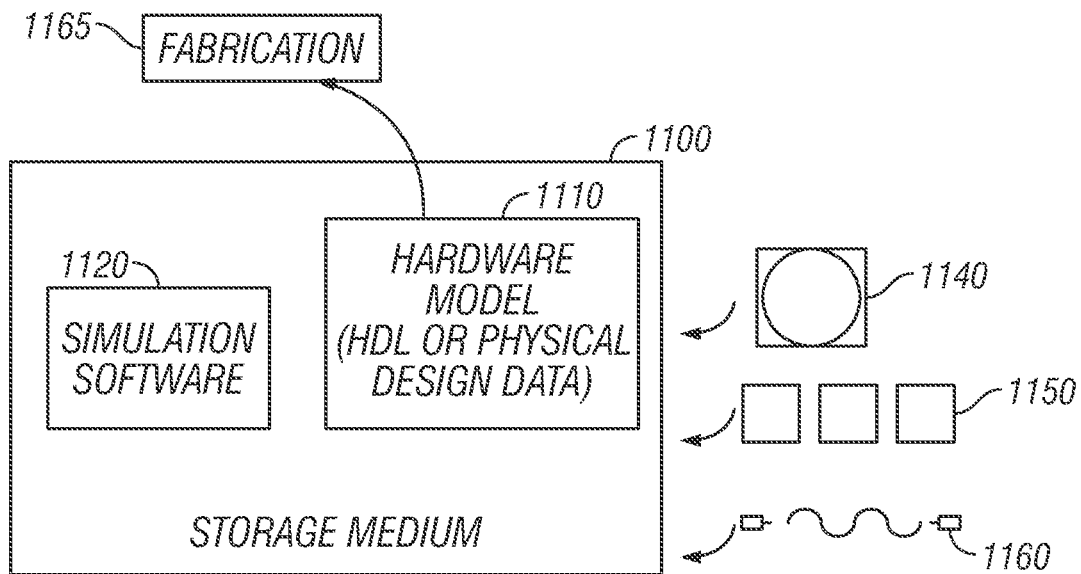


FIG. 11

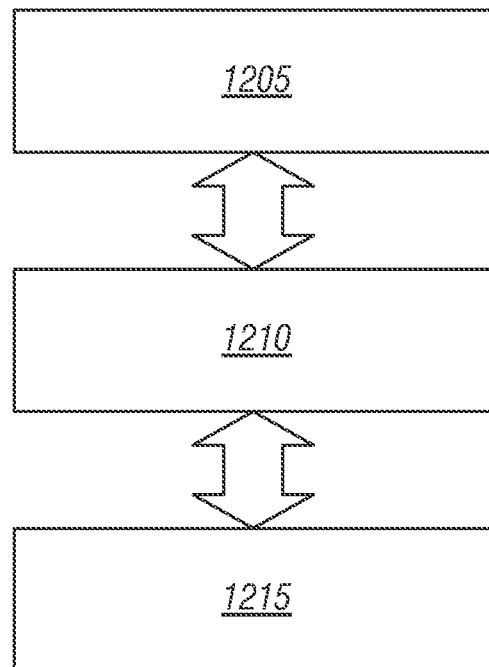


FIG. 12

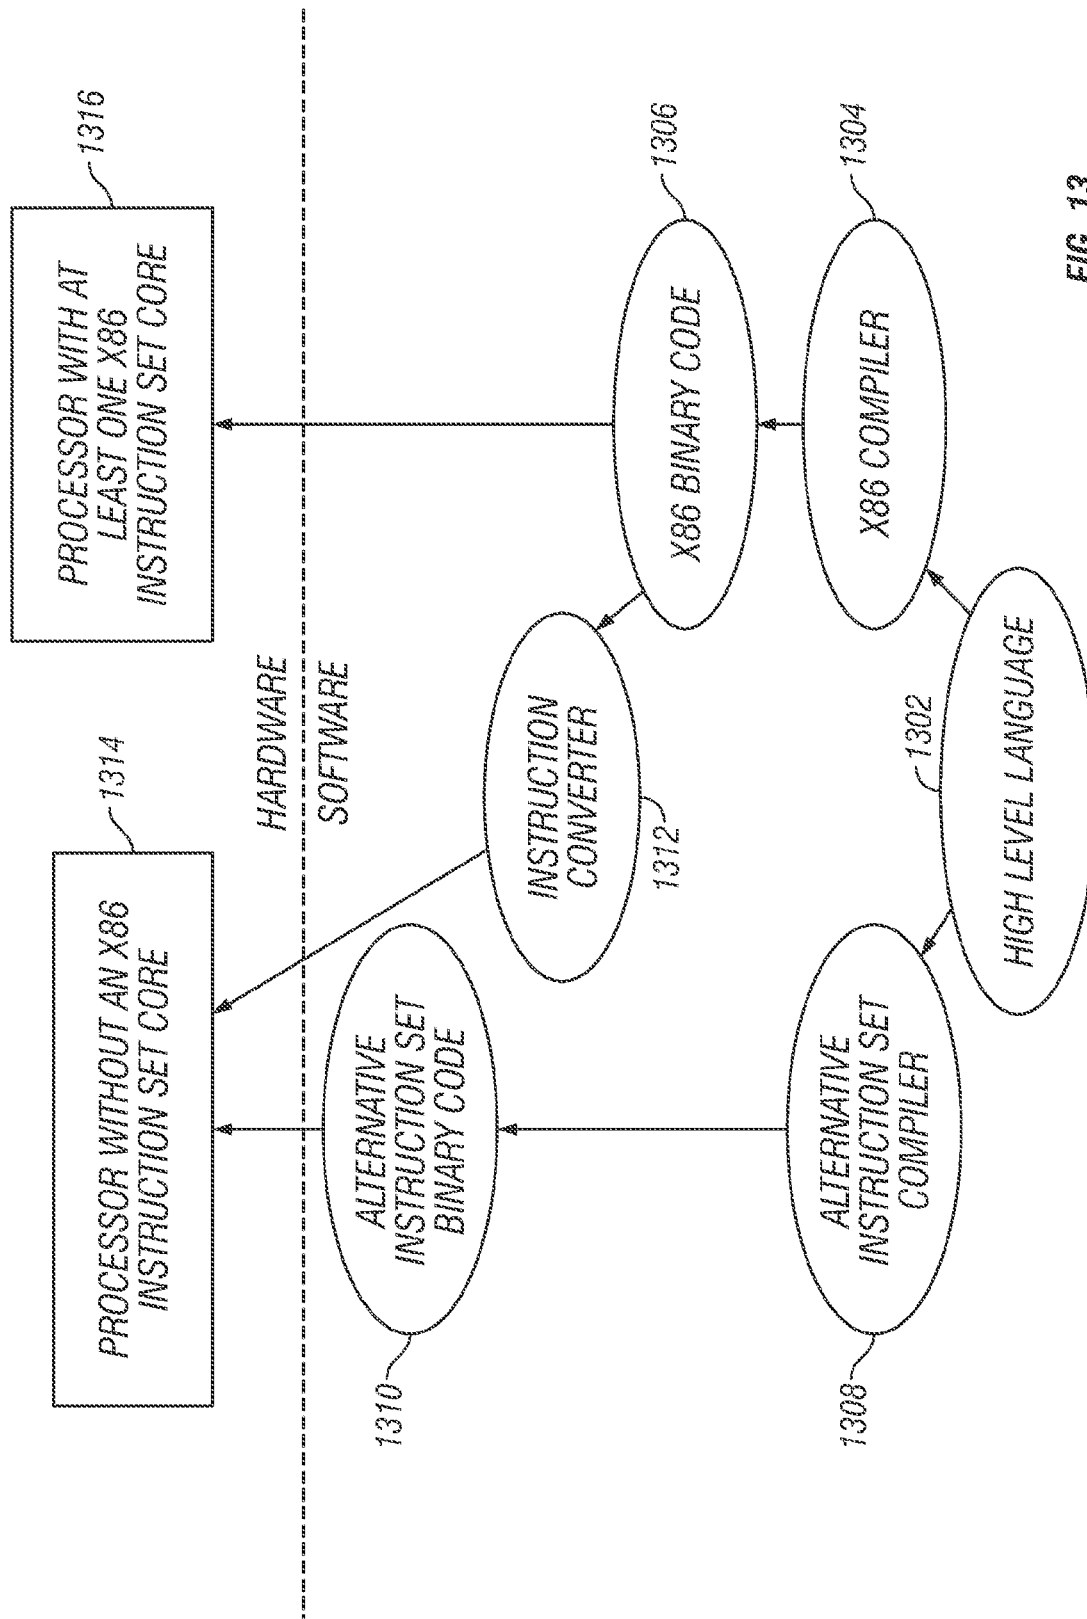


FIG. 13

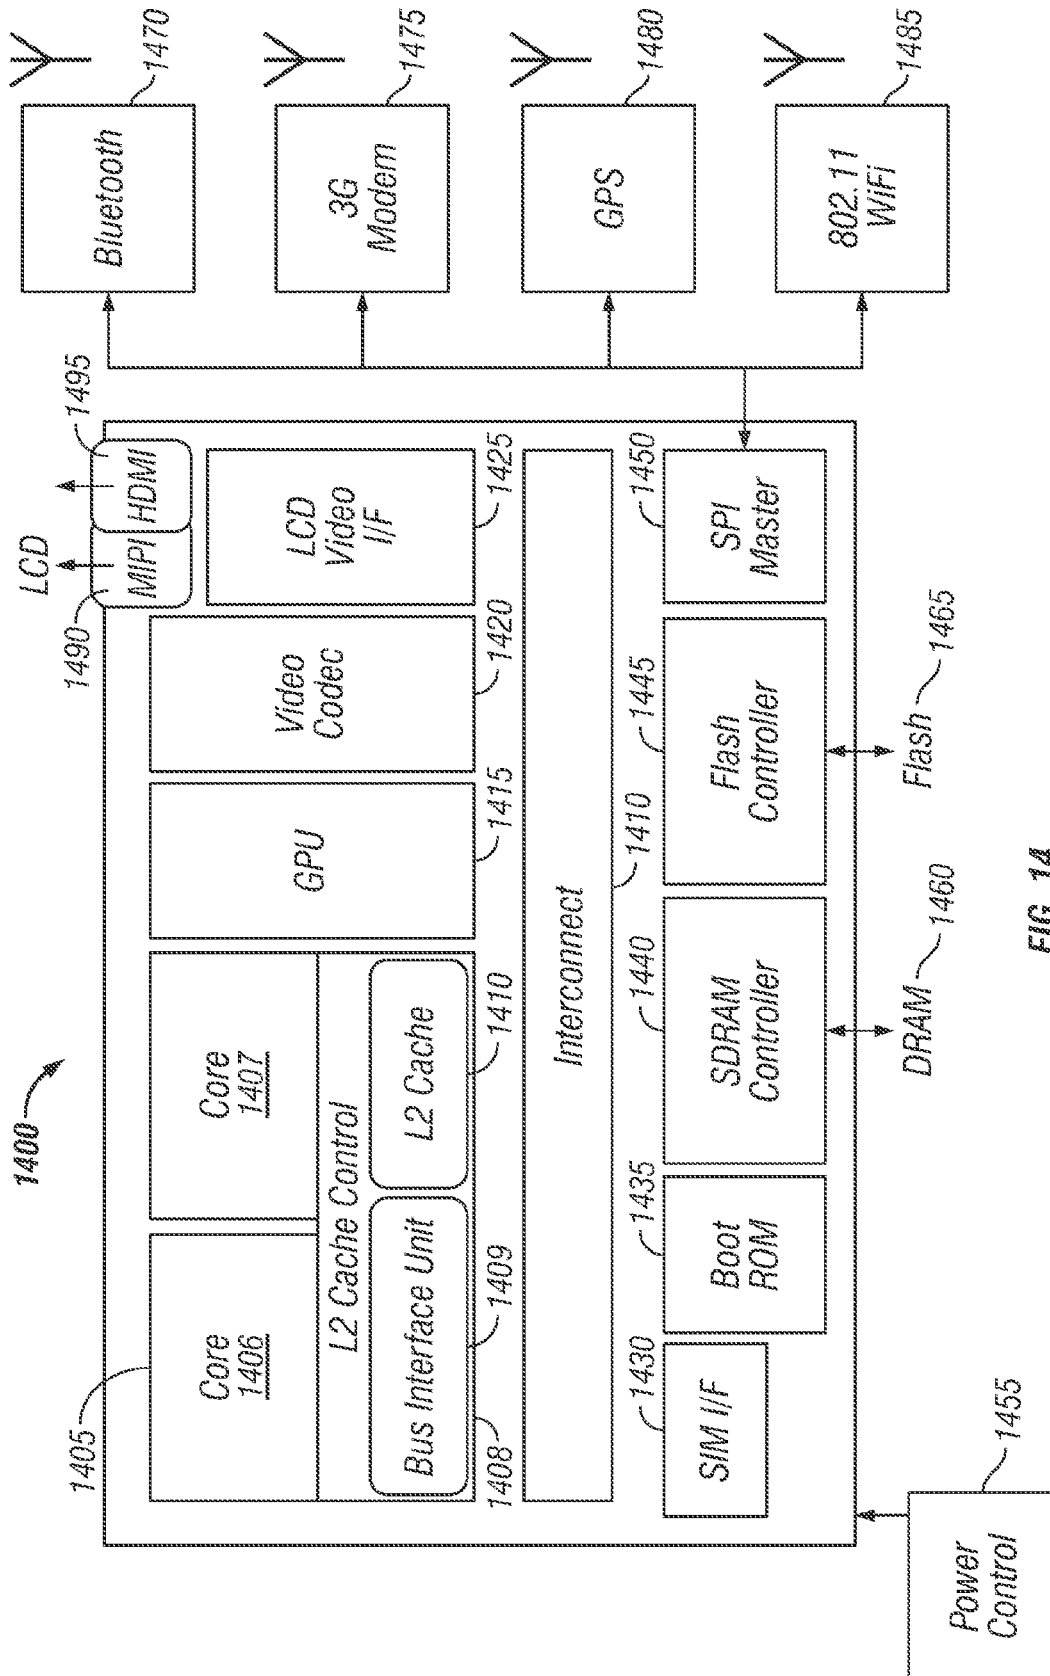


FIG. 14

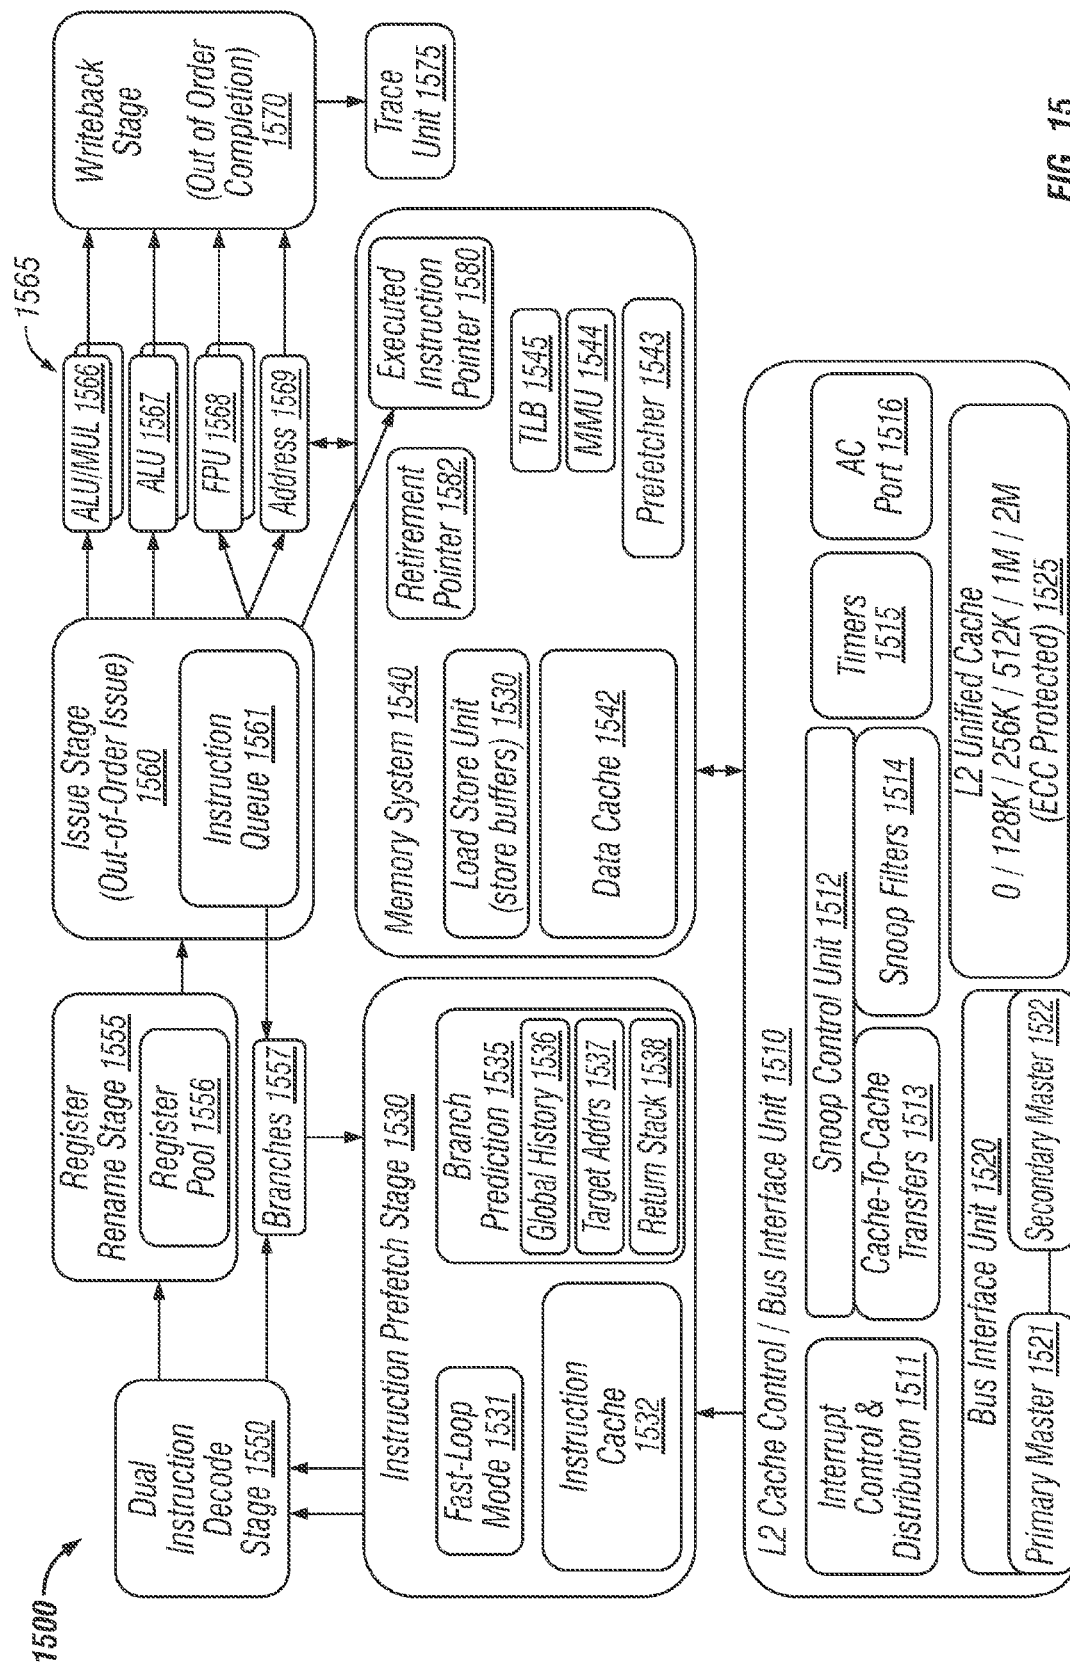


FIG. 15

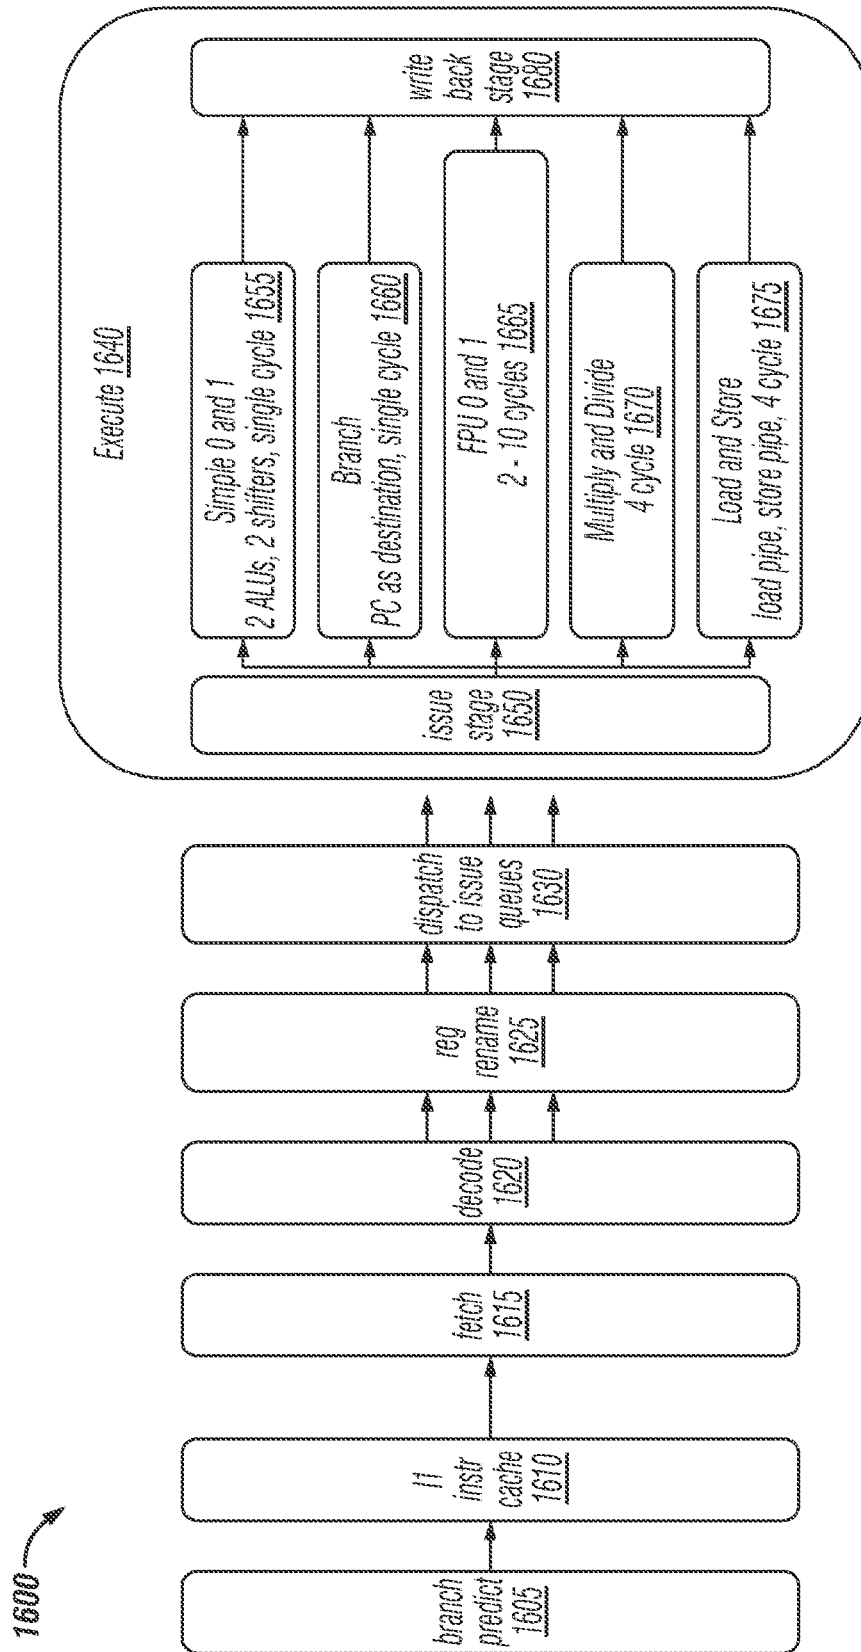


FIG. 16

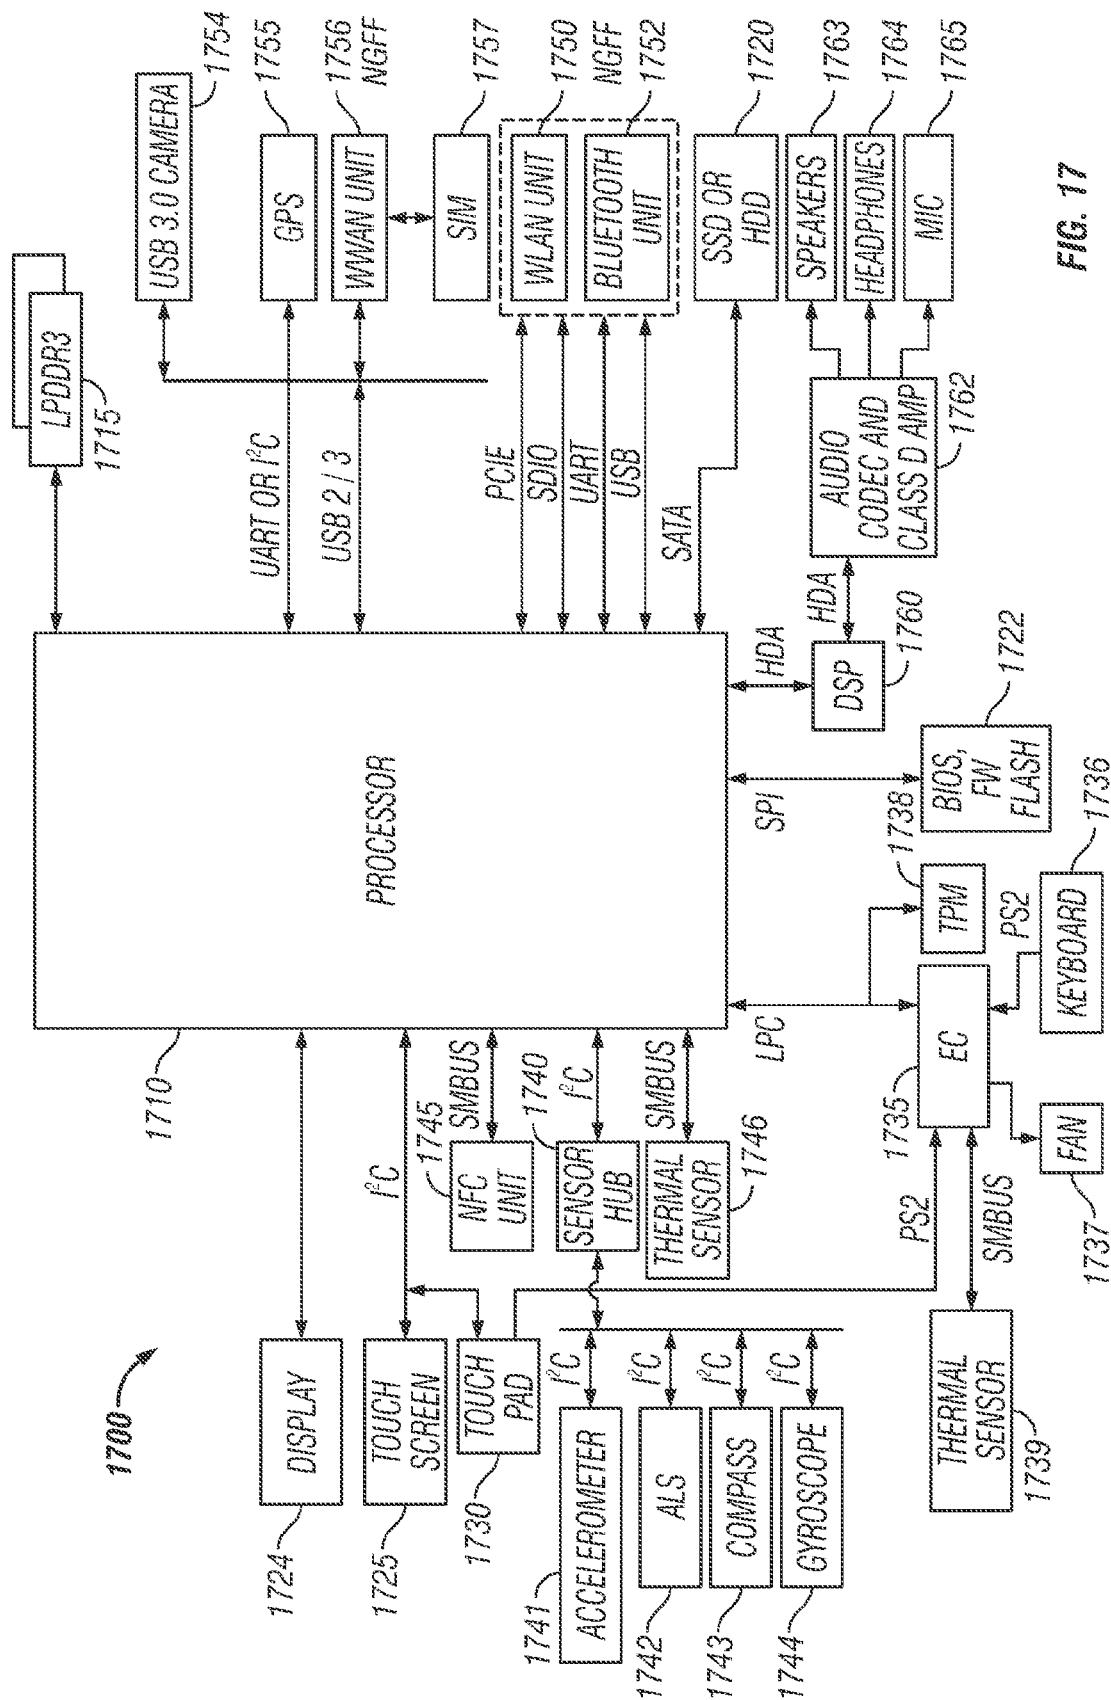
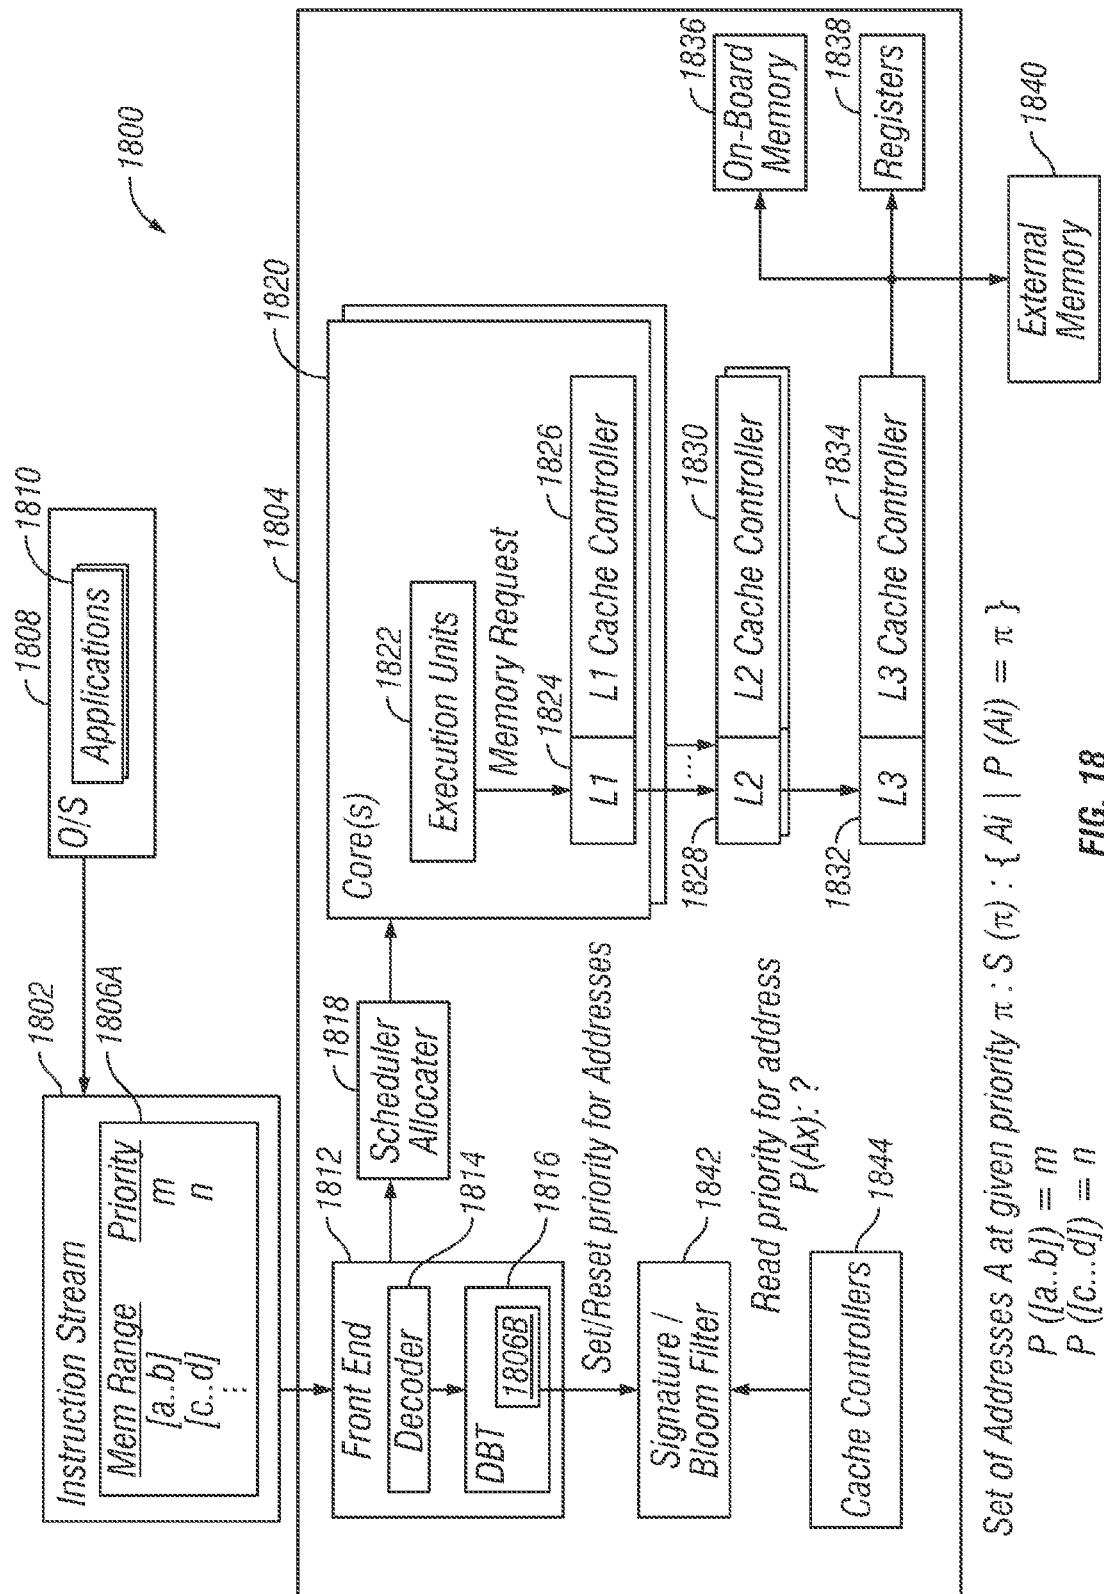


FIG. 17





924

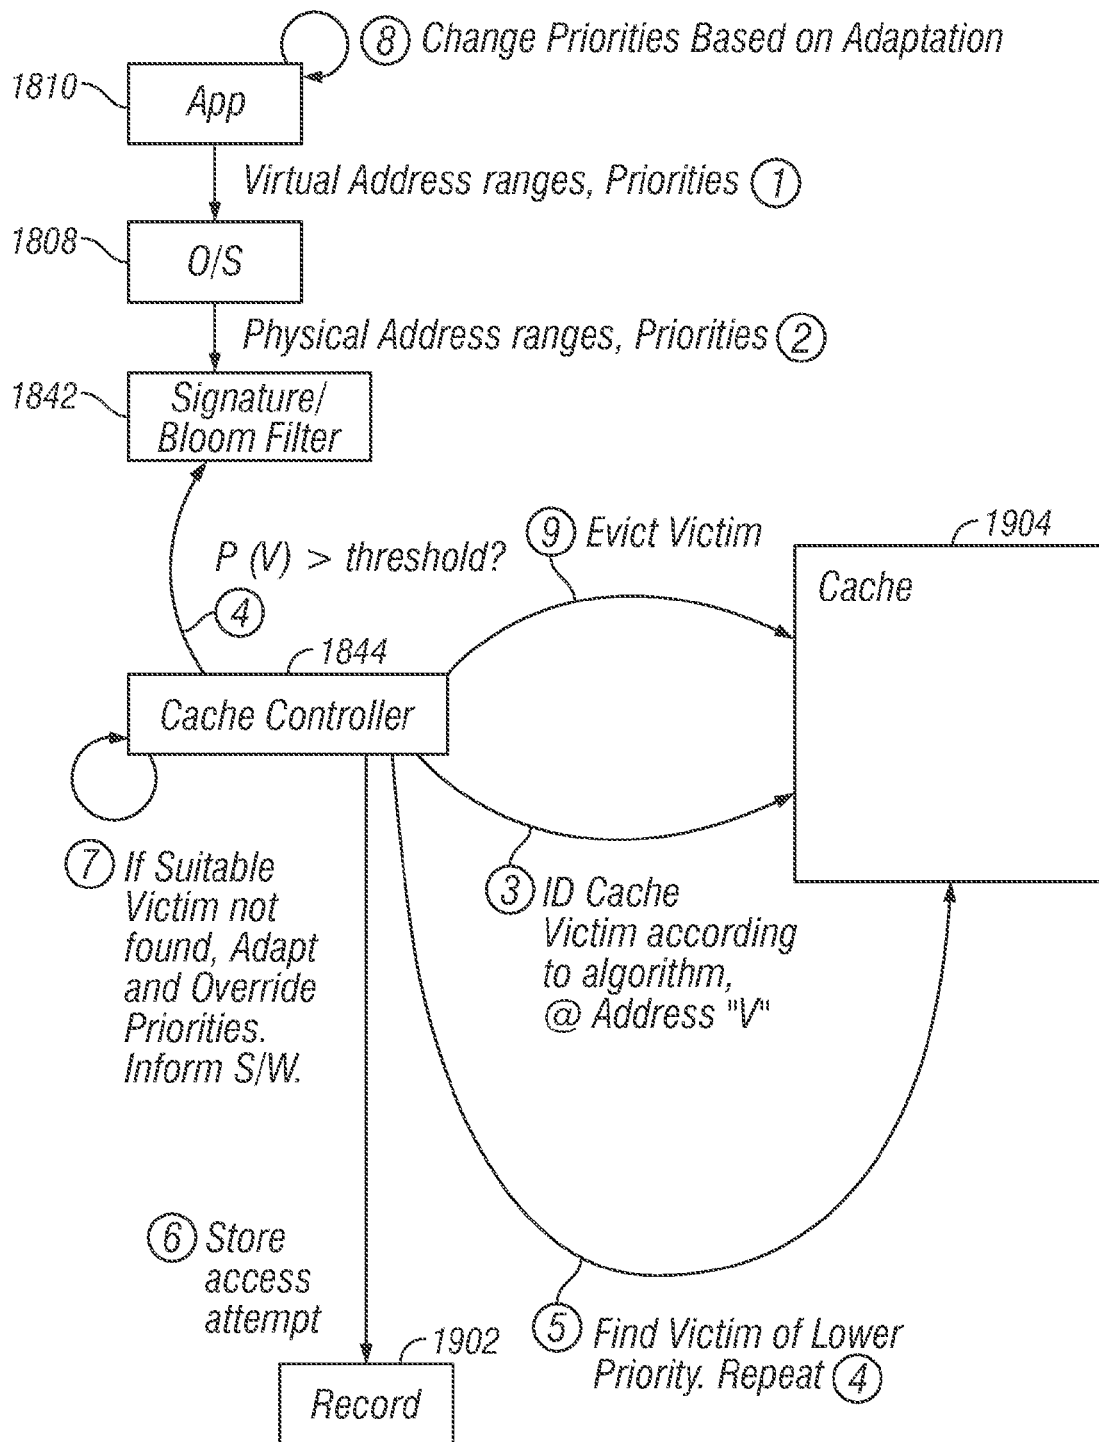
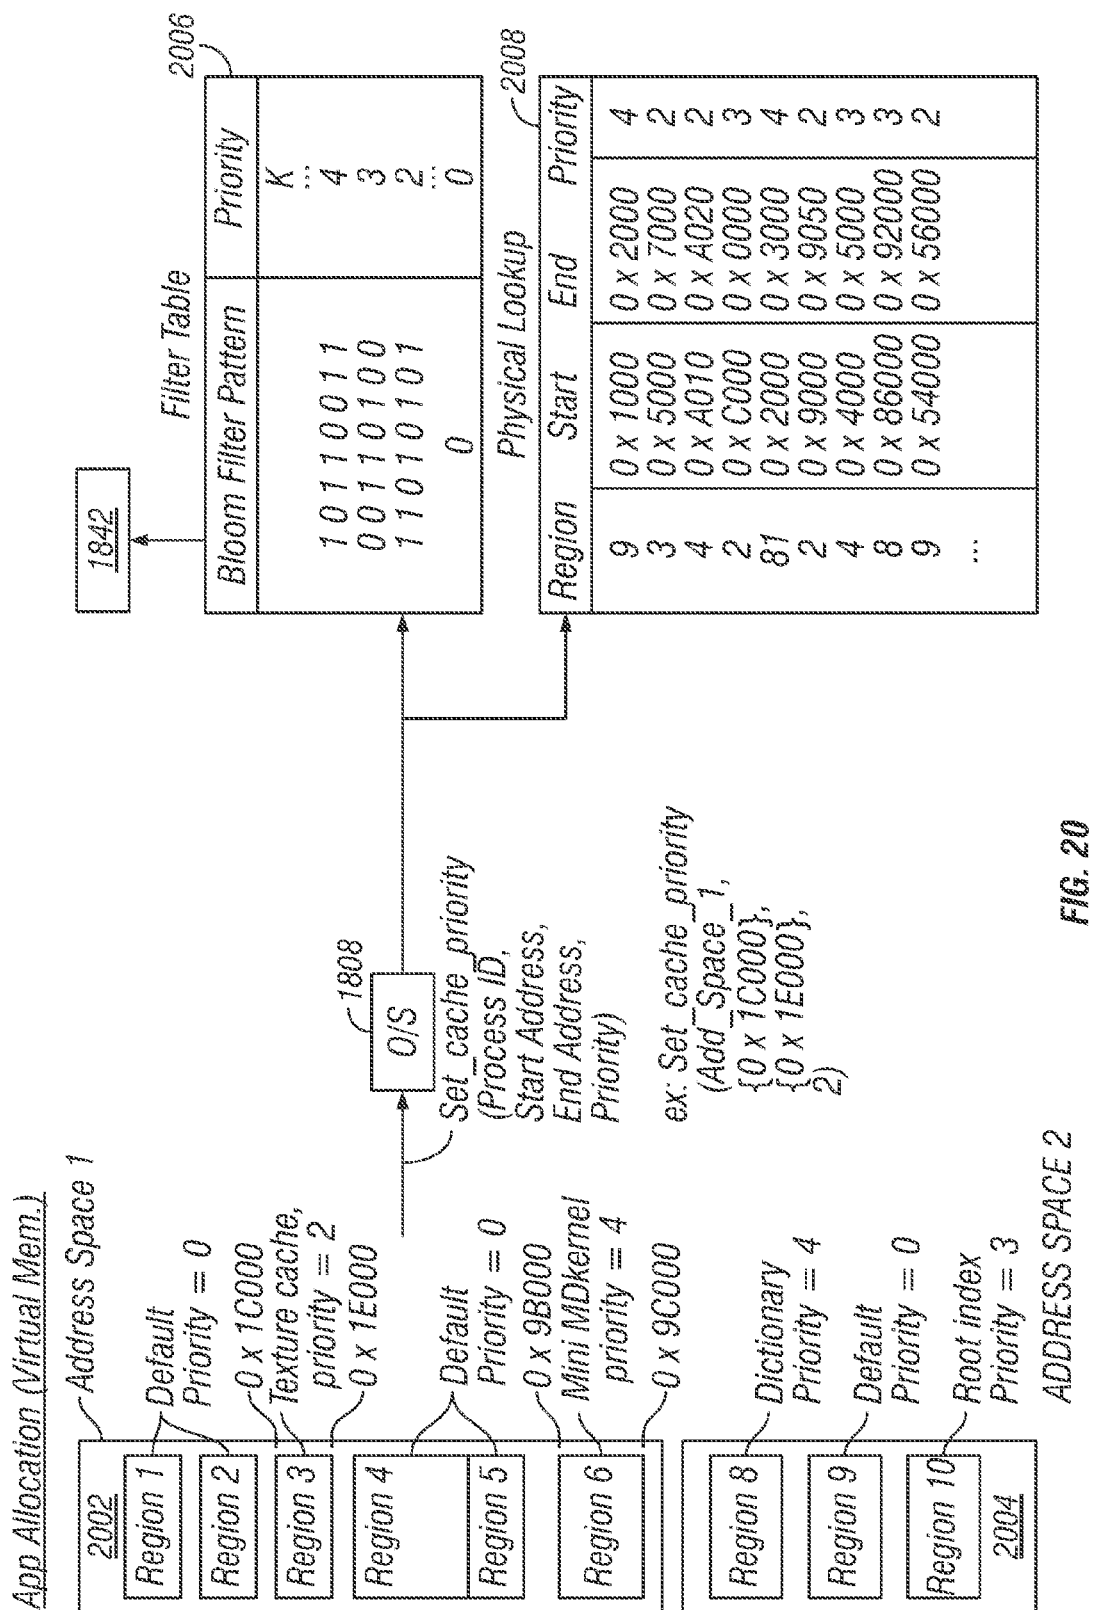


FIG. 19



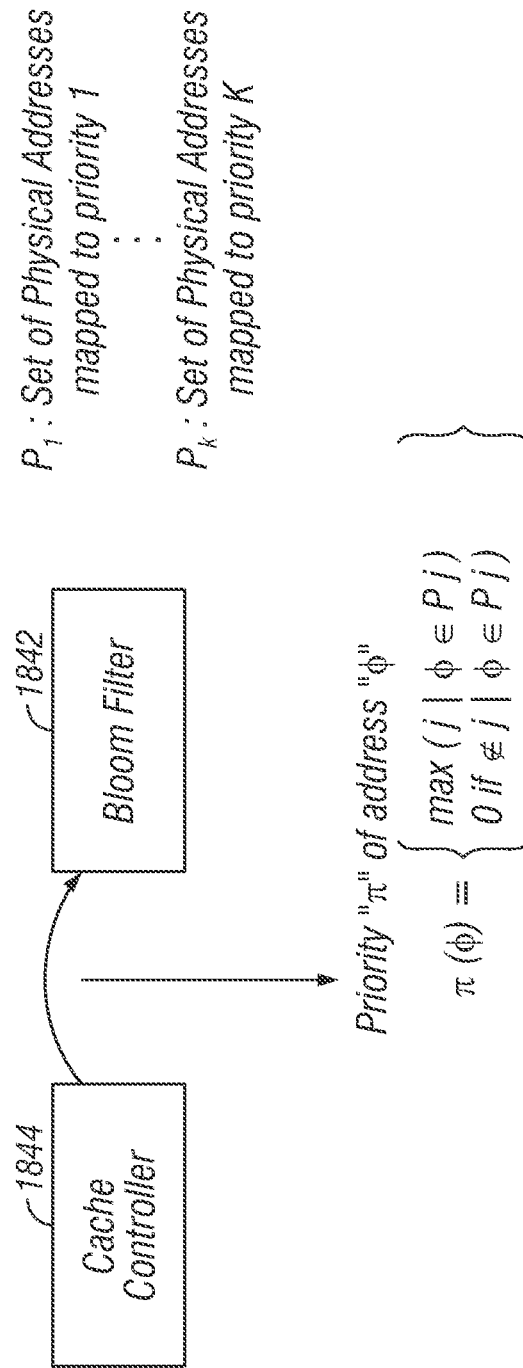


FIG. 20  
(Cont'd)

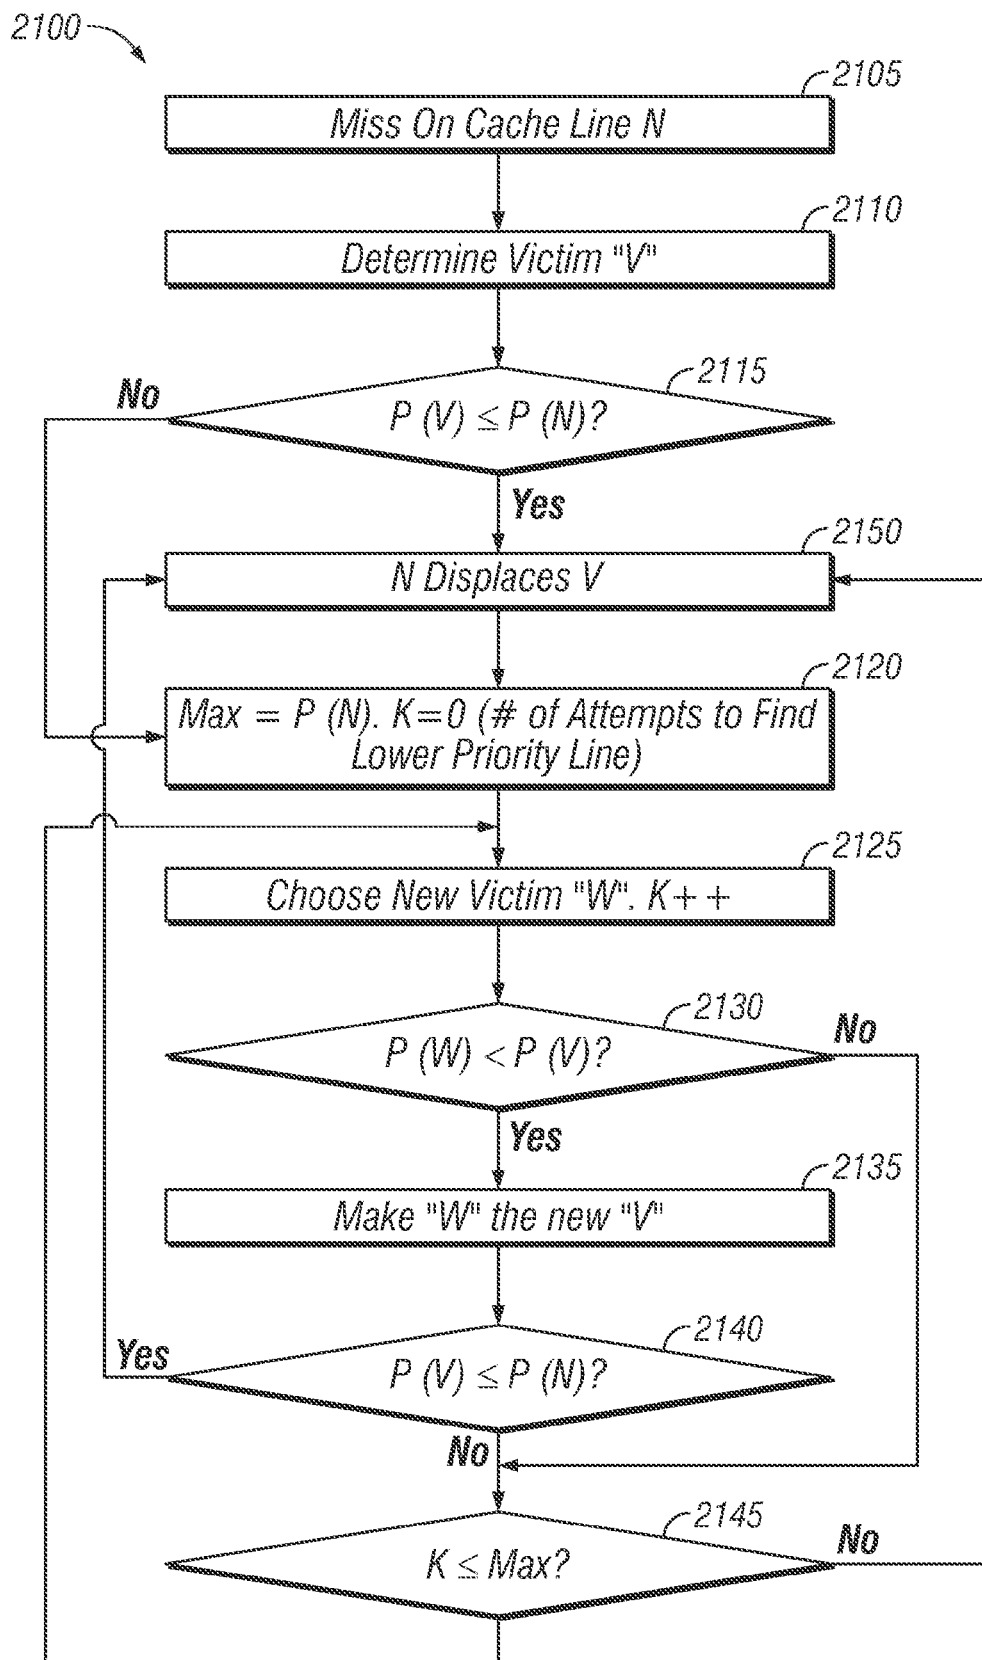


FIG. 21

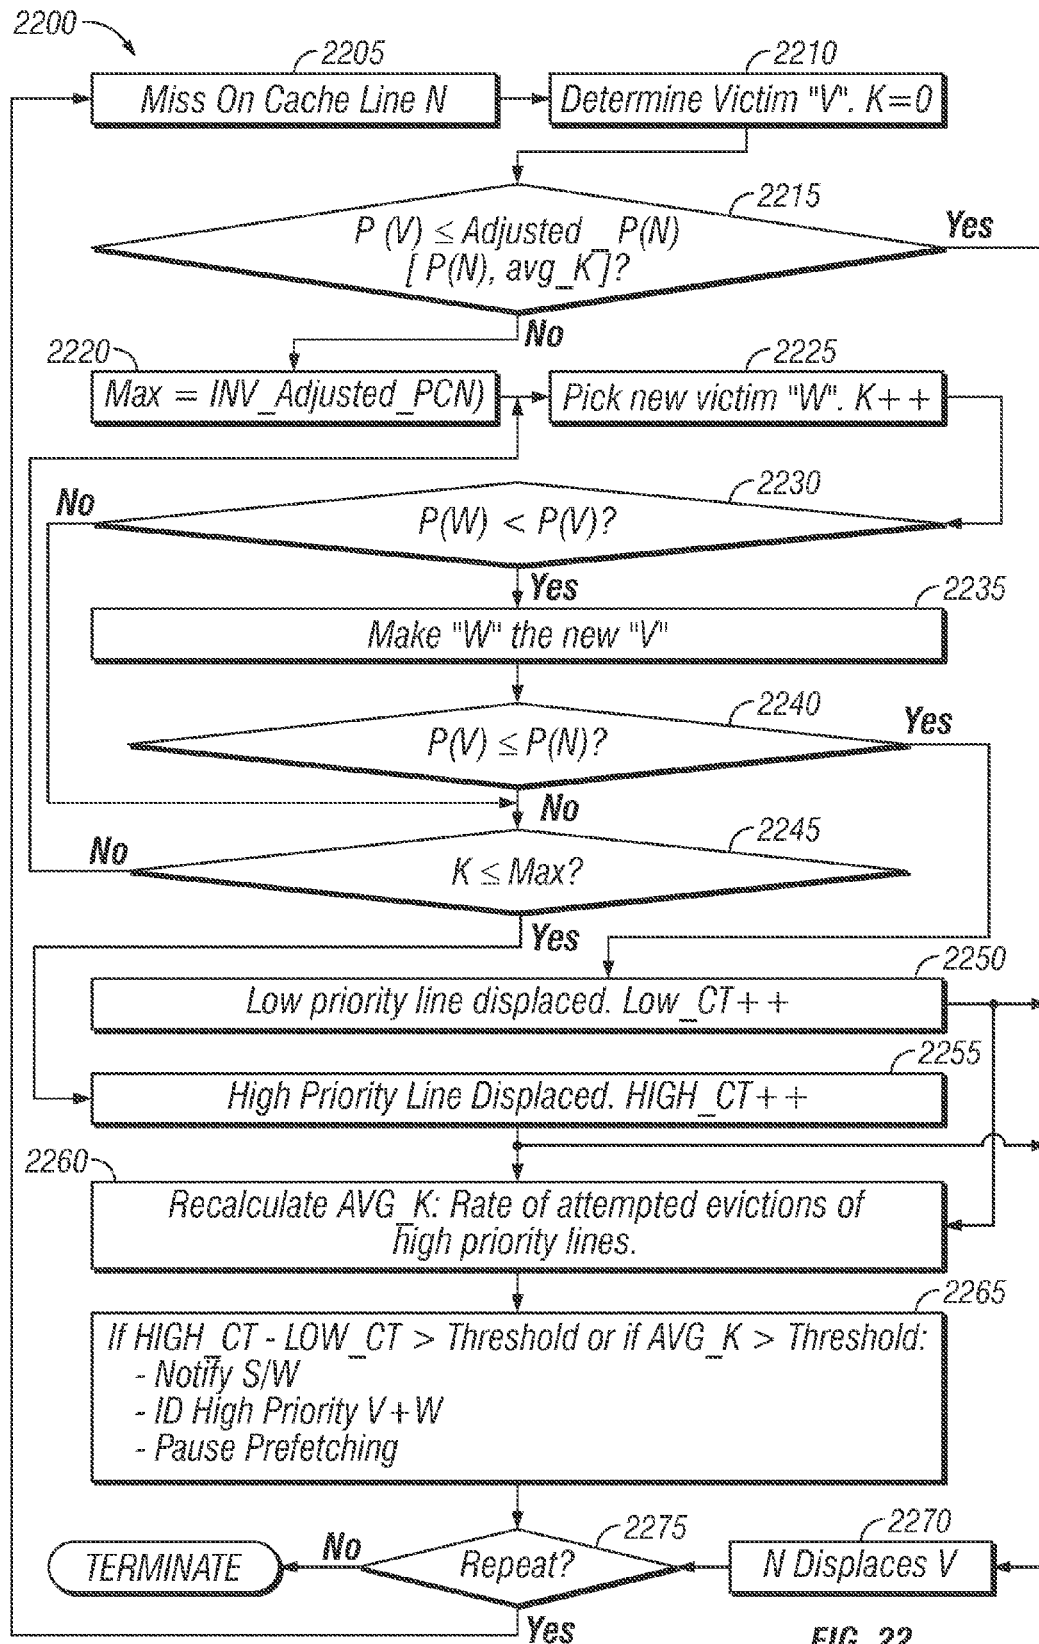


FIG. 22

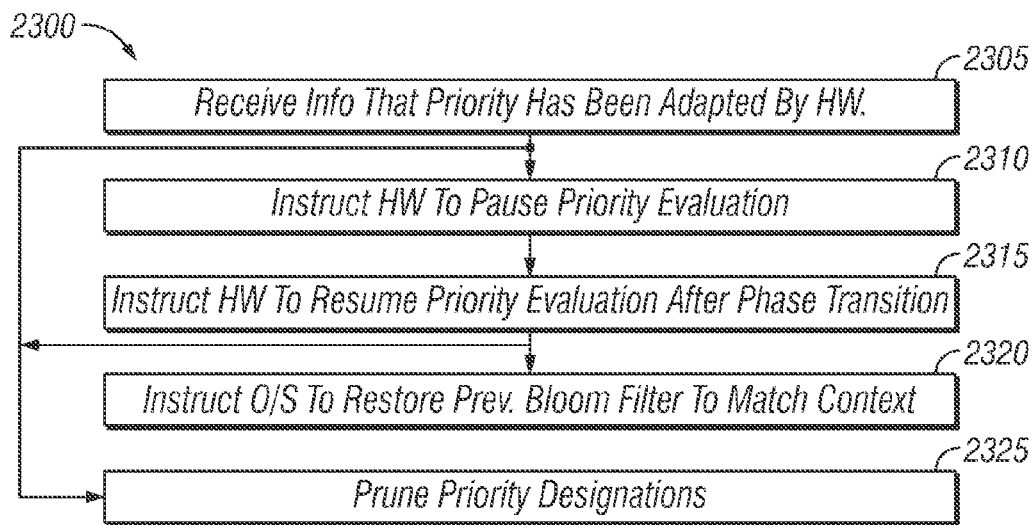


FIG. 23

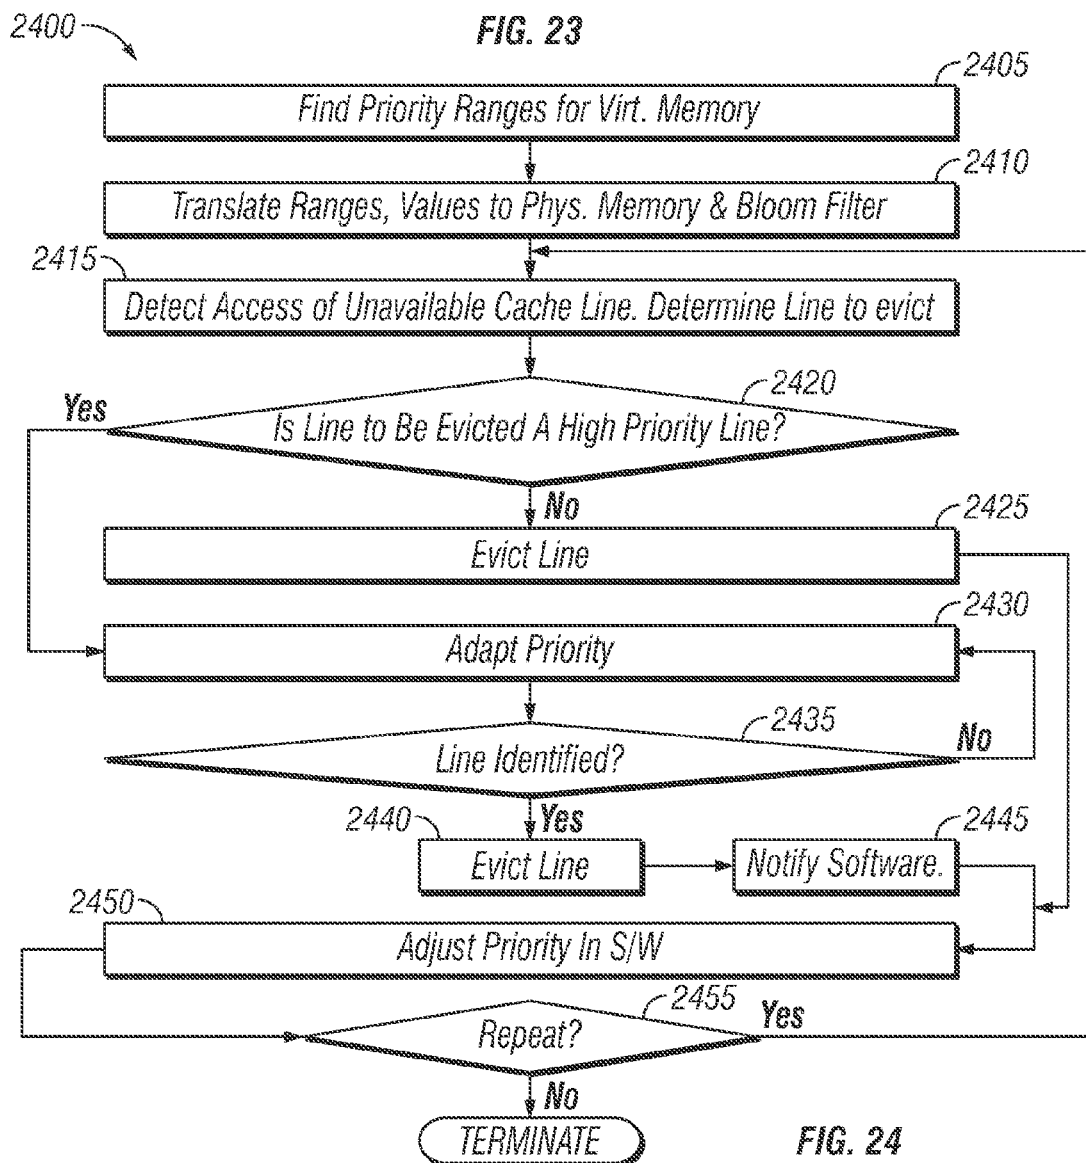


FIG. 24

# INSTRUCTION AND LOGIC FOR ADAPTIVE DATASET PRIORITIES IN PROCESSOR CACHES

## FIELD OF THE INVENTION

The present disclosure pertains to the field of processing logic, microprocessors, and associated instruction set architecture that, when executed by the processor or other processing logic, perform logical, mathematical, or other functional operations.

## DESCRIPTION OF RELATED ART

Multiprocessor systems are becoming more and more common. Applications of multiprocessor systems include dynamic domain partitioning all the way down to desktop computing. In order to take advantage of multiprocessor systems, code to be executed may be separated into multiple threads for execution by various processing entities. Each thread may be executed in parallel with one another. Furthermore, in order to increase the utility of a processing entity, out-of-order execution may be employed. Out-of-order execution may execute instructions as input to such instructions is made available. Thus, an instruction that appears later in a code sequence may be executed before an instruction appearing earlier in a code sequence.

## DESCRIPTION OF THE FIGURES

Embodiments are illustrated by way of example and not limitation in the Figures of the accompanying drawings:

FIG. 1A is a block diagram of an exemplary computer system formed with a processor that may include execution units to execute an instruction, in accordance with embodiments of the present disclosure;

FIG. 1B illustrates a data processing system, in accordance with embodiments of the present disclosure;

FIG. 1C illustrates other embodiments of a data processing system for performing single instruction multiple data operations;

FIG. 2 is a block diagram of the micro-architecture for a processor that may include logic circuits to perform instructions, in accordance with embodiments of the present disclosure;

FIG. 3A illustrates various packed data type representations in multimedia registers, in accordance with embodiments of the present disclosure;

FIG. 3B illustrates possible in-register data storage formats, in accordance with embodiments of the present disclosure;

FIG. 3C illustrates various signed and unsigned packed data type representations in multimedia registers, in accordance with embodiments of the present disclosure;

FIG. 3D illustrates an embodiment of an operation encoding format;

FIG. 3E illustrates another possible operation encoding format having forty or more bits, in accordance with embodiments of the present disclosure;

FIG. 3F illustrates yet another possible operation encoding format, in accordance with embodiments of the present disclosure;

FIG. 4A is a block diagram illustrating an in-order pipeline and a register renaming stage, out-of-order issue/execution pipeline, in accordance with embodiments of the present disclosure;

FIG. 4B is a block diagram illustrating an in-order architecture core and a register renaming logic, out-of-order issue/execution logic to be included in a processor, in accordance with embodiments of the present disclosure;

FIG. 5A is a block diagram of a processor, in accordance with embodiments of the present disclosure;

FIG. 5B is a block diagram of an example implementation of a core, in accordance with embodiments of the present disclosure;

FIG. 6 is a block diagram of a system, in accordance with embodiments of the present disclosure;

FIG. 7 is a block diagram of a second system, in accordance with embodiments of the present disclosure;

FIG. 8 is a block diagram of a third system in accordance with embodiments of the present disclosure;

FIG. 9 is a block diagram of a system-on-a-chip, in accordance with embodiments of the present disclosure;

FIG. 10 illustrates a processor containing a central processing unit and a graphics processing unit which may perform at least one instruction, in accordance with embodiments of the present disclosure;

FIG. 11 is a block diagram illustrating the development of IP cores, in accordance with embodiments of the present disclosure;

FIG. 12 illustrates how an instruction of a first type may be emulated by a processor of a different type, in accordance with embodiments of the present disclosure;

FIG. 13 illustrates a block diagram contrasting the use of a software instruction converter to convert binary instructions in a source instruction set to binary instructions in a target instruction set, in accordance with embodiments of the present disclosure;

FIG. 14 is a block diagram of an architecture of a processor, in accordance with embodiments of the present disclosure;

FIG. 15 is a more detailed block diagram of an architecture of a processor, in accordance with embodiments of the present disclosure;

FIG. 16 is a block diagram of an execution pipeline for an architecture of a processor, in accordance with embodiments of the present disclosure;

FIG. 17 is a block diagram of an electronic device for utilizing a processor, in accordance with embodiments of the present disclosure;

FIG. 18 is a block diagram of a system for implementing an instruction and logic for adaptive dataset priorities, in accordance with embodiments of the present disclosure;

FIG. 19 is an illustration of operation of system to perform evaluation and adaptation of priority datasets during a cache miss, in accordance with embodiments of the present disclosure;

FIG. 20 is an illustration of operation of system to specify dataset priorities, according to embodiments of the present disclosure;

FIG. 21 illustrates example operation of system to perform and adapt cache eviction, according to embodiments of the present disclosure;

FIG. 22 illustrates further example operation of system to perform and adapt cache eviction, according to embodiments of the present disclosure;

FIG. 23 is an illustration of operation of software to perform corrective action when notified by hardware that too many high-priority addresses have been evicted from cache, according to embodiments of the present disclosure; and

FIG. 24 is a flowchart of an example embodiment of a method for executing adaptive dataset priorities, in accordance with embodiments of the present disclosure.



The following description describes an instruction and processing logic for adaptive data set priorities with a processor, virtual processor, package, computer system, or other processing apparatus. Such adaptive dataset priorities may be used to evaluate contents of caches of such a processor. In the following description, numerous specific details such as processing logic, processor types, micro-architectural conditions, events, enablement mechanisms, and the like are set forth in order to provide a more thorough understanding of embodiments of the present disclosure. It will be appreciated, however, by one skilled in the art that the embodiments may be practiced without such specific details. Additionally, some well-known structures, circuits, and the like have not been shown in detail to avoid unnecessarily obscuring embodiments of the present disclosure.

Although the following embodiments are described with reference to a processor, other embodiments are applicable to other types of integrated circuits and logic devices. Similar techniques and teachings of embodiments of the present disclosure may be applied to other types of circuits or semiconductor devices that may benefit from higher pipeline throughput and improved performance. The teachings of embodiments of the present disclosure are applicable to any processor or machine that performs data manipulations. However, the embodiments are not limited to processors or machines that perform 512-bit, 256-bit, 128-bit, 64-bit, 32-bit, or 16-bit data operations and may be applied to any processor and machine in which manipulation or management of data may be performed. In addition, the following description provides examples, and the accompanying drawings show various examples for the purposes of illustration. However, these examples should not be construed in a limiting sense as they are merely intended to provide examples of embodiments of the present disclosure rather than to provide an exhaustive list of all possible implementations of embodiments of the present disclosure.

Although the below examples describe instruction handling and distribution in the context of execution units and logic circuits, other embodiments of the present disclosure may be accomplished by way of a data or instructions stored on a machine-readable, tangible medium, which when performed by a machine cause the machine to perform functions consistent with at least one embodiment of the disclosure. In one embodiment, functions associated with embodiments of the present disclosure are embodied in machine-executable instructions. The instructions may be used to cause a general-purpose or special-purpose processor that may be programmed with the instructions to perform the steps of the present disclosure. Embodiments of the present disclosure may be provided as a computer program product or software which may include a machine or computer-readable medium having stored thereon instructions which may be used to program a computer (or other electronic devices) to perform one or more operations according to embodiments of the present disclosure. Furthermore, steps of embodiments of the present disclosure might be performed by specific hardware components that contain fixed-function logic for performing the steps, or by any combination of programmed computer components and fixed-function hardware components.

Instructions used to program logic to perform embodiments of the present disclosure may be stored within a memory in the system, such as DRAM, cache, flash memory, or other storage. Furthermore, the instructions may be distributed via a network or by way of other computer-readable media. Thus a machine-readable medium may include any

mechanism for storing or transmitting information in a form readable by a machine (e.g., a computer), but is not limited to, floppy diskettes, optical disks, Compact Disc, Read-Only Memory (CD-ROMs), and magneto-optical disks, Read-Only Memory (ROMs), Random Access Memory (RAM), Erasable Programmable Read-Only Memory (EPROM), Electrically Erasable Programmable Read-Only Memory (EEPROM), magnetic or optical cards, flash memory, or a tangible, machine-readable storage used in the transmission of information over the Internet via electrical, optical, acoustic or other forms of propagated signals (e.g., carrier waves, infrared signals, digital signals, etc.). Accordingly, the computer-readable medium may include any type of tangible machine-readable medium suitable for storing or transmitting electronic instructions or information in a form readable by a machine (e.g., a computer).

A design may go through various stages, from creation to simulation to fabrication. Data representing a design may represent the design in a number of manners. First, as may be useful in simulations, the hardware may be represented using a hardware description language or another functional description language. Additionally, a circuit level model with logic and/or transistor gates may be produced at some stages of the design process. Furthermore, designs, at some stage, may reach a level of data representing the physical placement of various devices in the hardware model. In cases wherein some semiconductor fabrication techniques are used, the data representing the hardware model may be the data specifying the presence or absence of various features on different mask layers for masks used to produce the integrated circuit. In any representation of the design, the data may be stored in any form of a machine-readable medium. A memory or a magnetic or optical storage such as a disc may be the machine-readable medium to store information transmitted via optical or electrical wave modulated or otherwise generated to transmit such information. When an electrical carrier wave indicating or carrying the code or design is transmitted, to the extent that copying, buffering, or retransmission of the electrical signal is performed, a new copy may be made. Thus, a communication provider or a network provider may store on a tangible, machine-readable medium, at least temporarily, an article, such as information encoded into a carrier wave, embodying techniques of embodiments of the present disclosure.

In modern processors, a number of different execution units may be used to process and execute a variety of code and instructions. Some instructions may be quicker to complete while others may take a number of clock cycles to complete. The faster the throughput of instructions, the better the overall performance of the processor. Thus it would be advantageous to have as many instructions execute as fast as possible. However, there may be certain instructions that have greater complexity and require more in terms of execution time and processor resources, such as floating point instructions, load/store operations, data moves, etc.

As more computer systems are used in internet, text, and multimedia applications, additional processor support has been introduced over time. In one embodiment, an instruction set may be associated with one or more computer architectures, including data types, instructions, register architecture, addressing modes, memory architecture, interrupt and exception handling, and external input and output (I/O).

In one embodiment, the instruction set architecture (ISA) may be implemented by one or more micro-architectures, which may include processor logic and circuits used to implement one or more instruction sets. Accordingly, processors with different micro-architectures may share at least a portion

of a common instruction set. For example, Intel® Pentium 4 processors, Intel® Core™ processors, and processors from Advanced Micro Devices, Inc. of Sunnyvale Calif. implement nearly identical versions of the x86 instruction set (with some extensions that have been added with newer versions), but have different internal designs. Similarly, processors designed by other processor development companies, such as ARM Holdings, Ltd., MIPS, or their licensees or adopters, may share at least a portion a common instruction set, but may include different processor designs. For example, the same register architecture of the ISA may be implemented in different ways in different micro-architectures using new or well-known techniques, including dedicated physical registers, one or more dynamically allocated physical registers using a register renaming mechanism (e.g., the use of a Register Alias Table (RAT), a Reorder Buffer (ROB) and a retirement register file. In one embodiment, registers may include one or more registers, register architectures, register files, or other register sets that may or may not be addressable by a software programmer.

An instruction may include one or more instruction formats. In one embodiment, an instruction format may indicate various fields (number of bits, location of bits, etc.) to specify, among other things, the operation to be performed and the operands on which that operation will be performed. In a further embodiment, some instruction formats may be further defined by instruction templates (or sub-formats). For example, the instruction templates of a given instruction format may be defined to have different subsets of the instruction format's fields and/or defined to have a given field interpreted differently. In one embodiment, an instruction may be expressed using an instruction format (and, if defined, in a given one of the instruction templates of that instruction format) and specifies or indicates the operation and the operands upon which the operation will operate.

Scientific, financial, auto-vectorized general purpose, RMS (recognition, mining, and synthesis), and visual and multimedia applications (e.g., 2D/3D graphics, image processing, video compression/decompression, voice recognition algorithms and audio manipulation) may require the same operation to be performed on a large number of data items. In one embodiment, Single Instruction Multiple Data (SIMD) refers to a type of instruction that causes a processor to perform an operation on multiple data elements. SIMD technology may be used in processors that may logically divide the bits in a register into a number of fixed-sized or variable-sized data elements, each of which represents a separate value. For example, in one embodiment, the bits in a 64-bit register may be organized as a source operand containing four separate 16-bit data elements, each of which represents a separate 16-bit value. This type of data may be referred to as 'packed' data type or 'vector' data type, and operands of this data type may be referred to as packed data operands or vector operands. In one embodiment, a packed data item or vector may be a sequence of packed data elements stored within a single register, and a packed data operand or a vector operand may be a source or destination operand of a SIMD instruction (or 'packed data instruction' or a 'vector instruction'). In one embodiment, a SIMD instruction specifies a single vector operation to be performed on two source vector operands to generate a destination vector operand (also referred to as a result vector operand) of the same or different size, with the same or different number of data elements, and in the same or different data element order.

SIMD technology, such as that employed by the Intel® Core™ processors having an instruction set including x86, MMX™, Streaming SIMD Extensions (SSE), SSE2, SSE3,

SSE4.1, and SSE4.2 instructions, ARM processors, such as the ARM Cortex® family of processors having an instruction set including the Vector Floating Point (VFP) and/or NEON instructions, and MIPS processors, such as the Loongson family of processors developed by the Institute of Computing Technology (ICT) of the Chinese Academy of Sciences, has enabled a significant improvement in application performance (Core™ and MMX™ are registered trademarks or trademarks of Intel Corporation of Santa Clara, Calif.).

In one embodiment, destination and source registers/data may be generic terms to represent the source and destination of the corresponding data or operation. In some embodiments, they may be implemented by registers, memory, or other storage areas having other names or functions than those depicted. For example, in one embodiment, "DEST1" may be a temporary storage register or other storage area, whereas "SRC1" and "SRC2" may be a first and second source storage register or other storage area, and so forth. In other embodiments, two or more of the SRC and DEST storage areas may correspond to different data storage elements within the same storage area (e.g., a SIMD register). In one embodiment, one of the source registers may also act as a destination register by, for example, writing back the result of an operation performed on the first and second source data to one of the two source registers serving as a destination registers.

FIG. 1A is a block diagram of an exemplary computer system formed with a processor that may include execution units to execute an instruction, in accordance with embodiments of the present disclosure. System 100 may include a component, such as a processor 102 to employ execution units including logic to perform algorithms for process data, in accordance with the present disclosure, such as in the embodiment described herein. System 100 may be representative of processing systems based on the PENTIUM® III, PENTIUM® 4, Xeon™, Itanium®, XScale™ and/or StrongARM™ microprocessors available from Intel Corporation of Santa Clara, Calif., although other systems (including PCs having other microprocessors, engineering workstations, set-top boxes and the like) may also be used. In one embodiment, sample system 100 may execute a version of the WINDOWS™ operating system available from Microsoft Corporation of Redmond, Wash., although other operating systems (UNIX and Linux for example), embedded software, and/or graphical user interfaces, may also be used. Thus, embodiments of the present disclosure are not limited to any specific combination of hardware circuitry and software.

Embodiments are not limited to computer systems. Embodiments of the present disclosure may be used in other devices such as handheld devices and embedded applications. Some examples of handheld devices include cellular phones, Internet Protocol devices, digital cameras, personal digital assistants (PDAs), and handheld PCs. Embedded applications may include a micro controller, a digital signal processor (DSP), system on a chip, network computers (NetPC), set-top boxes, network hubs, wide area network (WAN) switches, or any other system that may perform one or more instructions in accordance with at least one embodiment.

Computer system 100 may include a processor 102 that may include one or more execution units 108 to perform an algorithm to perform at least one instruction in accordance with one embodiment of the present disclosure. One embodiment may be described in the context of a single processor desktop or server system, but other embodiments may be included in a multiprocessor system. System 100 may be an example of a 'hub' system architecture. System 100 may include a processor 102 for processing data signals. Processor

**102** may include a complex instruction set computer (CISC) microprocessor, a reduced instruction set computing (RISC) microprocessor, a very long instruction word (VLIW) microprocessor, a processor implementing a combination of instruction sets, or any other processor device, such as a digital signal processor, for example. In one embodiment, processor **102** may be coupled to a processor bus **110** that may transmit data signals between processor **102** and other components in system **100**. The elements of system **100** may perform conventional functions that are well known to those familiar with the art.

In one embodiment, processor **102** may include a Level 1 (L1) internal cache memory **104**. Depending on the architecture, the processor **102** may have a single internal cache or multiple levels of internal cache. In another embodiment, the cache memory may reside external to processor **102**. Other embodiments may also include a combination of both internal and external caches depending on the particular implementation and needs. Register file **106** may store different types of data in various registers including integer registers, floating point registers, status registers, and instruction pointer register.

Execution unit **108**, including logic to perform integer and floating point operations, also resides in processor **102**. Processor **102** may also include a microcode (ucode) ROM that stores microcode for certain macroinstructions. In one embodiment, execution unit **108** may include logic to handle a packed instruction set **109**. By including the packed instruction set **109** in the instruction set of a general-purpose processor **102**, along with associated circuitry to execute the instructions, the operations used by many multimedia applications may be performed using packed data in a general-purpose processor **102**. Thus, many multimedia applications may be accelerated and executed more efficiently by using the full width of a processor's data bus for performing operations on packed data. This may eliminate the need to transfer smaller units of data across the processor's data bus to perform one or more operations one data element at a time.

Embodiments of an execution unit **108** may also be used in micro controllers, embedded processors, graphics devices, DSPs, and other types of logic circuits. System **100** may include a memory **120**. Memory **120** may be implemented as a dynamic random access memory (DRAM) device, a static random access memory (SRAM) device, flash memory device, or other memory device. Memory **120** may store instructions and/or data represented by data signals that may be executed by processor **102**.

A system logic chip **116** may be coupled to processor bus **110** and memory **120**. System logic chip **116** may include a memory controller hub (MCH). Processor **102** may communicate with MCH **116** via a processor bus **110**. MCH **116** may provide a high bandwidth memory path **118** to memory **120** for instruction and data storage and for storage of graphics commands, data and textures. MCH **116** may direct data signals between processor **102**, memory **120**, and other components in system **100** and to bridge the data signals between processor bus **110**, memory **120**, and system I/O **122**. In some embodiments, the system logic chip **116** may provide a graphics port for coupling to a graphics controller **112**. MCH **116** may be coupled to memory **120** through a memory interface **118**. Graphics card **112** may be coupled to MCH **116** through an Accelerated Graphics Port (AGP) interconnect **114**.

System **100** may use a proprietary hub interface bus **122** to couple MCH **116** to I/O controller hub (ICH) **130**. In one embodiment, ICH **130** may provide direct connections to some I/O devices via a local I/O bus. The local I/O bus may

include a high-speed I/O bus for connecting peripherals to memory **120**, chipset, and processor **102**. Examples may include the audio controller, firmware hub (flash BIOS) **128**, wireless transceiver **126**, data storage **124**, legacy I/O controller containing user input and keyboard interfaces, a serial expansion port such as Universal Serial Bus (USB), and a network controller **134**. Data storage device **124** may comprise a hard disk drive, a floppy disk drive, a CD-ROM device, a flash memory device, or other mass storage device.

For another embodiment of a system, an instruction in accordance with one embodiment may be used with a system on a chip. One embodiment of a system on a chip comprises of a processor and a memory. The memory for one such system may include a flash memory. The flash memory may be located on the same die as the processor and other system components. Additionally, other logic blocks such as a memory controller or graphics controller may also be located on a system on a chip.

FIG. 1B illustrates a data processing system **140** which implements the principles of embodiments of the present disclosure. It will be readily appreciated by one of skill in the art that the embodiments described herein may operate with alternative processing systems without departure from the scope of embodiments of the disclosure.

Computer system **140** comprises a processing core **159** for performing at least one instruction in accordance with one embodiment. In one embodiment, processing core **159** represents a processing unit of any type of architecture, including but not limited to a CISC, a RISC or a VLIW type architecture. Processing core **159** may also be suitable for manufacture in one or more process technologies and by being represented on a machine-readable media in sufficient detail, may be suitable to facilitate said manufacture.

Processing core **159** comprises an execution unit **142**, a set of register files **145**, and a decoder **144**. Processing core **159** may also include additional circuitry (not shown) which may be unnecessary to the understanding of embodiments of the present disclosure. Execution unit **142** may execute instructions received by processing core **159**. In addition to performing typical processor instructions, execution unit **142** may perform instructions in packed instruction set **143** for performing operations on packed data formats. Packed instruction set **143** may include instructions for performing embodiments of the disclosure and other packed instructions. Execution unit **142** may be coupled to register file **145** by an internal bus. Register file **145** may represent a storage area on processing core **159** for storing information, including data. As previously mentioned, it is understood that the storage area may store the packed data might not be critical. Execution unit **142** may be coupled to decoder **144**. Decoder **144** may decode instructions received by processing core **159** into control signals and/or microcode entry points. In response to these control signals and/or microcode entry points, execution unit **142** performs the appropriate operations. In one embodiment, the decoder may interpret the opcode of the instruction, which will indicate what operation should be performed on the corresponding data indicated within the instruction.

Processing core **159** may be coupled with bus **141** for communicating with various other system devices, which may include but are not limited to, for example, synchronous dynamic random access memory (SDRAM) control **146**, static random access memory (SRAM) control **147**, burst flash memory interface **148**, personal computer memory card international association (PCMCIA)/compact flash (CF) card control **149**, liquid crystal display (LCD) control **150**, direct memory access (DMA) controller **151**, and alternative bus

master interface **152**. In one embodiment, data processing system **140** may also comprise an I/O bridge **154** for communicating with various I/O devices via an I/O bus **153**. Such I/O devices may include but are not limited to, for example, universal asynchronous receiver/transmitter (UART) **155**, universal serial bus (USB) **156**, Bluetooth wireless UART **157** and I/O expansion interface **158**.

One embodiment of data processing system **140** provides for mobile, network and/or wireless communications and a processing core **159** that may perform SIMD operations including a text string comparison operation. Processing core **159** may be programmed with various audio, video, imaging and communications algorithms including discrete transformations such as a Walsh-Hadamard transform, a fast Fourier transform (FFT), a discrete cosine transform (DCT), and their respective inverse transforms; compression/decompression techniques such as color space transformation, video encode motion estimation or video decode motion compensation; and modulation/demodulation (MODEM) functions such as pulse coded modulation (PCM).

FIG. 1C illustrates other embodiments of a data processing system that performs SIMD operations. In one embodiment, data processing system **160** may include a main processor **166**, a SIMD coprocessor **161**, a cache memory **167**, and an input/output system **168**. Input/output system **168** may optionally be coupled to a wireless interface **169**. SIMD coprocessor **161** may perform operations including instructions in accordance with one embodiment. In one embodiment, processing core **170** may be suitable for manufacture in one or more process technologies and by being represented on a machine-readable media in sufficient detail, may be suitable to facilitate the manufacture of all or part of data processing system **160** including processing core **170**.

In one embodiment, SIMD coprocessor **161** comprises an execution unit **162** and a set of register files **164**. One embodiment of main processor **165** comprises a decoder **165** to recognize instructions of instruction set **163** including instructions in accordance with one embodiment for execution by execution unit **162**. In other embodiments, SIMD coprocessor **161** also comprises at least part of decoder **165** to decode instructions of instruction set **163**. Processing core **170** may also include additional circuitry (not shown) which may be unnecessary to the understanding of embodiments of the present disclosure.

In operation, main processor **166** executes a stream of data processing instructions that control data processing operations of a general type including interactions with cache memory **167**, and input/output system **168**. Embedded within the stream of data processing instructions may be SIMD coprocessor instructions. Decoder **165** of main processor **166** recognizes these SIMD coprocessor instructions as being of a type that should be executed by an attached SIMD coprocessor **161**. Accordingly, main processor **166** issues these SIMD coprocessor instructions (or control signals representing SIMD coprocessor instructions) on the coprocessor bus **166**. From coprocessor bus **166**, these instructions may be received by any attached SIMD coprocessors. In this case, SIMD coprocessor **161** may accept and execute any received SIMD coprocessor instructions intended for it.

Data may be received via wireless interface **169** for processing by the SIMD coprocessor instructions. For one example, voice communication may be received in the form of a digital signal, which may be processed by the SIMD coprocessor instructions to regenerate digital audio samples representative of the voice communications. For another example, compressed audio and/or video may be received in the form of a digital bit stream, which may be processed by

the SIMD coprocessor instructions to regenerate digital audio samples and/or motion video frames. In one embodiment of processing core **170**, main processor **166**, and a SIMD coprocessor **161** may be integrated into a single processing core **170** comprising an execution unit **162**, a set of register files **164**, and a decoder **165** to recognize instructions of instruction set **163** including instructions in accordance with one embodiment.

FIG. 2 is a block diagram of the micro-architecture for a processor **200** that may include logic circuits to perform instructions, in accordance with embodiments of the present disclosure. In some embodiments, an instruction in accordance with one embodiment may be implemented to operate on data elements having sizes of byte, word, doubleword, quadword, etc., as well as datatypes, such as single and double precision integer and floating point datatypes. In one embodiment, in-order front end **201** may implement a part of processor **200** that may fetch instructions to be executed and prepares the instructions to be used later in the processor pipeline. Front end **201** may include several units. In one embodiment, instruction prefetcher **226** fetches instructions from memory and feeds the instructions to an instruction decoder **228** which in turn decodes or interprets the instructions. For example, in one embodiment, the decoder decodes a received instruction into one or more operations called “micro-instructions” or “micro-operations” (also called micro op or uops) that the machine may execute. In other embodiments, the decoder parses the instruction into an opcode and corresponding data and control fields that may be used by the micro-architecture to perform operations in accordance with one embodiment. In one embodiment, trace cache **230** may assemble decoded uops into program ordered sequences or traces in uop queue **234** for execution. When trace cache **230** encounters a complex instruction, microcode ROM **232** provides the uops needed to complete the operation.

Some instructions may be converted into a single micro-op, whereas others need several micro-ops to complete the full operation. In one embodiment, if more than four micro-ops are needed to complete an instruction, decoder **228** may access microcode ROM **232** to perform the instruction. In one embodiment, an instruction may be decoded into a small number of micro ops for processing at instruction decoder **228**. In another embodiment, an instruction may be stored within microcode ROM **232** should a number of micro-ops be needed to accomplish the operation. Trace cache **230** refers to an entry point programmable logic array (PLA) to determine a correct micro-instruction pointer for reading the micro-code sequences to complete one or more instructions in accordance with one embodiment from micro-code ROM **232**. After microcode ROM **232** finishes sequencing micro-ops for an instruction, front end **201** of the machine may resume fetching micro-ops from trace cache **230**.

Out-of-order execution engine **203** may prepare instructions for execution. The out-of-order execution logic has a number of buffers to smooth out and re-order the flow of instructions to optimize performance as they go down the pipeline and get scheduled for execution. The allocator logic allocates the machine buffers and resources that each uop needs in order to execute. The register renaming logic renames logic registers onto entries in a register file. The allocator also allocates an entry for each uop in one of the two uop queues, one for memory operations and one for non-memory operations, in front of the instruction schedulers: memory scheduler, fast scheduler **202**, slow/general floating point scheduler **204**, and simple floating point scheduler **206**. Uop schedulers **202**, **204**, **206**, determine when a uop is ready

## 11

to execute based on the readiness of their dependent input register operand sources and the availability of the execution resources the uops need to complete their operation. Fast scheduler **202** of one embodiment may schedule on each half of the main clock cycle while the other schedulers may only schedule once per main processor clock cycle. The schedulers arbitrate for the dispatch ports to schedule uops for execution.

Register files **208**, **210** may be arranged between schedulers **202**, **204**, **206**, and execution units **212**, **214**, **216**, **218**, **220**, **222**, **224** in execution block **211**. Each of register files **208**, **210** perform integer and floating point operations, respectively. Each register file **208**, **210**, may include a bypass network that may bypass or forward just completed results that have not yet been written into the register file to new dependent uops. Integer register file **208** and floating point register file **210** may communicate data with the other. In one embodiment, integer register file **208** may be split into two separate register files, one register file for low-order thirty-two bits of data and a second register file for high order thirty-two bits of data. Floating point register file **210** may include 128-bit wide entries because floating point instructions typically have operands from 64 to 128 bits in width.

Execution block **211** may contain execution units **212**, **214**, **216**, **218**, **220**, **222**, **224**. Execution units **212**, **214**, **216**, **218**, **220**, **222**, **224** may execute the instructions. Execution block **211** may include register files **208**, **210** that store the integer and floating point data operand values that the micro-instructions need to execute. In one embodiment, processor **200** may comprise a number of execution units: address generation unit (AGU) **212**, AGU **214**, fast ALU **216**, fast ALU **218**, slow ALU **220**, floating point ALU **222**, floating point move unit **224**. In another embodiment, floating point execution blocks **222**, **224**, may execute floating point, MMX, SIMD, and SSE, or other operations. In yet another embodiment, floating point ALU **222** may include a 64-bit by 64-bit floating point divider to execute divide, square root, and remainder micro-ops. In various embodiments, instructions involving a floating point value may be handled with the floating point hardware. In one embodiment, ALU operations may be passed to high-speed ALU execution units **216**, **218**. High-speed ALUs **216**, **218** may execute fast operations with an effective latency of half a clock cycle. In one embodiment, most complex integer operations go to slow ALU **220** as slow ALU **220** may include integer execution hardware for long-latency type of operations, such as a multiplier, shifts, flag logic, and branch processing. Memory load/store operations may be executed by AGUs **212**, **214**. In one embodiment, integer ALUs **216**, **218**, **220** may perform integer operations on 64-bit data operands. In other embodiments, ALUs **216**, **218**, **220** may be implemented to support a variety of data bit sizes including sixteen, thirty-two, 128, 256, etc. Similarly, floating point units **222**, **224** may be implemented to support a range of operands having bits of various widths. In one embodiment, floating point units **222**, **224**, may operate on 128-bit wide packed data operands in conjunction with SIMD and multimedia instructions.

In one embodiment, uops schedulers **202**, **204**, **206**, dispatch dependent operations before the parent load has finished executing. As uops may be speculatively scheduled and executed in processor **200**, processor **200** may also include logic to handle memory misses. If a data load misses in the data cache, there may be dependent operations in flight in the pipeline that have left the scheduler with temporarily incorrect data. A replay mechanism tracks and re-executes instructions that use incorrect data. Only the dependent operations might need to be replayed and the independent ones may be allowed to complete. The schedulers and replay mechanism

## 12

of one embodiment of a processor may also be designed to catch instruction sequences for text string comparison operations.

The term “registers” may refer to the on-board processor storage locations that may be used as part of instructions to identify operands. In other words, registers may be those that may be usable from the outside of the processor (from a programmer’s perspective). However, in some embodiments registers might not be limited to a particular type of circuit. Rather, a register may store data, provide data, and perform the functions described herein. The registers described herein may be implemented by circuitry within a processor using any number of different techniques, such as dedicated physical registers, dynamically allocated physical registers using register renaming, combinations of dedicated and dynamically allocated physical registers, etc. In one embodiment, integer registers store 32-bit integer data. A register file of one embodiment also contains eight multimedia SIMD registers for packed data. For the discussions below, the registers may be understood to be data registers designed to hold packed data, such as 64-bit wide MMX™ registers (also referred to as ‘mm’ registers in some instances) in microprocessors enabled with MMX technology from Intel Corporation of Santa Clara, Calif. These MMX registers, available in both integer and floating point forms, may operate with packed data elements that accompany SIMD and SSE instructions. Similarly, 128-bit wide XMM registers relating to SSE2, SSE3, SSE4, or beyond (referred to generically as “SSEx”) technology may hold such packed data operands. In one embodiment, in storing packed data and integer data, the registers do not need to differentiate between the two data types. In one embodiment, integer and floating point may be contained in the same register file or different register files. Furthermore, in one embodiment, floating point and integer data may be stored in different registers or the same registers.

In the examples of the following figures, a number of data operands may be described. FIG. 3A illustrates various packed data type representations in multimedia registers, in accordance with embodiments of the present disclosure. FIG. 3A illustrates data types for a packed byte **310**, a packed word **320**, and a packed doubleword (dword) **330** for 128-bit wide operands. Packed byte format **310** of this example may be 128 bits long and contains sixteen packed byte data elements. A byte may be defined, for example, as eight bits of data. Information for each byte data element may be stored in bit 7 through bit 0 for byte 0, bit 15 through bit 8 for byte 1, bit 23 through bit 16 for byte 2, and finally bit 120 through bit 127 for byte 15. Thus, all available bits may be used in the register. This storage arrangement increases the storage efficiency of the processor. As well, with sixteen data elements accessed, one operation may now be performed on sixteen data elements in parallel.

Generally, a data element may include an individual piece of data that is stored in a single register or memory location with other data elements of the same length. In packed data sequences relating to SSEx technology, the number of data elements stored in a XMM register may be 128 bits divided by the length in bits of an individual data element. Similarly, in packed data sequences relating to MMX and SSE technology, the number of data elements stored in an MMX register may be 64 bits divided by the length in bits of an individual data element. Although the data types illustrated in FIG. 3A may be 128 bits long, embodiments of the present disclosure may also operate with 64-bit wide or other sized operands. Packed word format **320** of this example may be 128 bits long and contains eight packed word data elements. Each packed word contains sixteen bits of information. Packed doubleword for-

13

mat 330 of FIG. 3A may be 128 bits long and contains four packed doubleword data elements. Each packed doubleword data element contains thirty-two bits of information. A packed quadword may be 128 bits long and contain two packed quad-word data elements.

FIG. 3B illustrates possible in-register data storage formats, in accordance with embodiments of the present disclosure. Each packed data may include more than one independent data element. Three packed data formats are illustrated; packed half 341, packed single 342, and packed double 343. One embodiment of packed half 341, packed single 342, and packed double 343 contain fixed-point data elements. For another embodiment one or more of packed half 341, packed single 342, and packed double 343 may contain floating-point data elements. One embodiment of packed half 341 may be 128 bits long containing eight 16-bit data elements. One embodiment of packed single 342 may be 128 bits long and contains four 32-bit data elements. One embodiment of packed double 343 may be 128 bits long and contains two 64-bit data elements. It will be appreciated that such packed data formats may be further extended to other register lengths, for example, to 96-bits, 160-bits, 192-bits, 224-bits, 256-bits or more.

FIG. 3C illustrates various signed and unsigned packed data type representations in multimedia registers, in accordance with embodiments of the present disclosure. Unsigned packed byte representation 344 illustrates the storage of an unsigned packed byte in a SIMD register. Information for each byte data element may be stored in bit 7 through bit 0 for byte 0, bit 15 through bit 8 for byte 1, bit 23 through bit 16 for byte 2, and finally bit 120 through bit 127 for byte 15. Thus, all available bits may be used in the register. This storage arrangement may increase the storage efficiency of the processor. As well, with sixteen data elements accessed, one operation may now be performed on sixteen data elements in a parallel fashion. Signed packed byte representation 345 illustrates the storage of a signed packed byte. Note that the eighth bit of every byte data element may be the sign indicator. Unsigned packed word representation 346 illustrates how word seven through word zero may be stored in a SIMD register. Signed packed word representation 347 may be similar to the unsigned packed word in-register representation 346. Note that the sixteenth bit of each word data element may be the sign indicator. Unsigned packed doubleword representation 348 shows how doubleword data elements are stored. Signed packed doubleword representation 349 may be similar to unsigned packed doubleword in-register representation 348. Note that the necessary sign bit may be the thirty-second bit of each doubleword data element.

FIG. 3D illustrates an embodiment of an operation encoding (opcode). Furthermore, format 360 may include register/memory operand addressing modes corresponding with a type of opcode format described in the "IA-32 Intel Architecture Software Developer's Manual Volume 2: Instruction Set Reference," which is available from Intel Corporation, Santa Clara, Calif. on the world-wide-web (www) at intel.com/design/litcentr. In one embodiment, and instruction may be encoded by one or more of fields 361 and 362. Up to two operand locations per instruction may be identified, including up to two source operand identifiers 364 and 365. In one embodiment, destination operand identifier 366 may be the same as source operand identifier 364, whereas in other embodiments they may be different. In another embodiment, destination operand identifier 366 may be the same as source operand identifier 365, whereas in other embodiments they may be different. In one embodiment, one of the source operands identified by source operand identifiers 364 and 365

14

may be overwritten by the results of the text string comparison operations, whereas in other embodiments identifier 364 corresponds to a source register element and identifier 365 corresponds to a destination register element. In one embodiment, operand identifiers 364 and 365 may identify 32-bit or 64-bit source and destination operands.

FIG. 3E illustrates another possible operation encoding (opcode) format 370, having forty or more bits, in accordance with embodiments of the present disclosure. Opcode format 370 corresponds with opcode format 360 and comprises an optional prefix byte 378. An instruction according to one embodiment may be encoded by one or more of fields 378, 371, and 372. Up to two operand locations per instruction may be identified by source operand identifiers 374 and 375 and by prefix byte 378. In one embodiment, prefix byte 378 may be used to identify 32-bit or 64-bit source and destination operands. In one embodiment, destination operand identifier 376 may be the same as source operand identifier 374, whereas in other embodiments they may be different. For another embodiment, destination operand identifier 376 may be the same as source operand identifier 375, whereas in other embodiments they may be different. In one embodiment, an instruction operates on one or more of the operands identified by operand identifiers 374 and 375 and one or more operands identified by operand identifiers 374 and 375 may be overwritten by the results of the instruction, whereas in other embodiments, operands identified by identifiers 374 and 375 may be written to another data element in another register. Opcode formats 360 and 370 allow register to register, memory to register, register by memory, register by register, register by immediate, register to memory addressing specified in part by MOD fields 363 and 373 and by optional scale-index-base and displacement bytes.

FIG. 3F illustrates yet another possible operation encoding (opcode) format, in accordance with embodiments of the present disclosure. 64-bit single instruction multiple data (SIMD) arithmetic operations may be performed through a coprocessor data processing (CDP) instruction. Operation encoding (opcode) format 380 depicts one such CDP instruction having CDP opcode fields 382 and 0064 389. The type of CDP instruction, for another embodiment, operations may be encoded by one or more of fields 383, 384, 387, and 388. Up to three operand locations per instruction may be identified, including up to two source operand identifiers 385 and 390 and one destination operand identifier 386. One embodiment of the coprocessor may operate on eight, sixteen, thirty-two, and 64-bit values. In one embodiment, an instruction may be performed on integer data elements. In some embodiments, an instruction may be executed conditionally, using condition field 381. For some embodiments, source data sizes may be encoded by field 383. In some embodiments, Zero (Z), negative (N), carry (C), and overflow (V) detection may be done on SIMD fields. For some instructions, the type of saturation may be encoded by field 384.

FIG. 4A is a block diagram illustrating an in-order pipeline and a register renaming stage, out-of-order issue/execution pipeline, in accordance with embodiments of the present disclosure. FIG. 4B is a block diagram illustrating an in-order architecture core and a register renaming logic, out-of-order issue/execution logic to be included in a processor, in accordance with embodiments of the present disclosure. The solid lined boxes in FIG. 4A illustrate the in-order pipeline, while the dashed lined boxes illustrates the register renaming, out-of-order issue/execution pipeline. Similarly, the solid lined boxes in FIG. 4B illustrate the in-order architecture logic, while the dashed lined boxes illustrates the register renaming logic and out-of-order issue/execution logic.

## 15

In FIG. 4A, a processor pipeline 400 may include a fetch stage 402, a length decode stage 404, a decode stage 406, an allocation stage 408, a renaming stage 410, a scheduling (also known as a dispatch or issue) stage 412, a register read/memory read stage 414, an execute stage 416, a write-back/memory-write stage 418, an exception handling stage 422, and a commit stage 424.

In FIG. 4B, arrows denote a coupling between two or more units and the direction of the arrow indicates a direction of data flow between those units. FIG. 4B shows processor core 490 including a front end unit 430 coupled to an execution engine unit 450, and both may be coupled to a memory unit 470.

Core 490 may be a reduced instruction set computing (RISC) core, a complex instruction set computing (CISC) core, a very long instruction word (VLIW) core, or a hybrid or alternative core type. In one embodiment, core 490 may be a special-purpose core, such as, for example, a network or communication core, compression engine, graphics core, or the like.

Front end unit 430 may include a branch prediction unit 432 coupled to an instruction cache unit 434. Instruction cache unit 434 may be coupled to an instruction translation lookaside buffer (TLB) 436. TLB 436 may be coupled to an instruction fetch unit 438, which is coupled to a decode unit 440. Decode unit 440 may decode instructions, and generate as an output one or more micro-operations, micro-code entry points, microinstructions, other instructions, or other control signals, which may be decoded from, or which otherwise reflect, or may be derived from, the original instructions. The decoder may be implemented using various different mechanisms. Examples of suitable mechanisms include, but are not limited to, look-up tables, hardware implementations, programmable logic arrays (PLAs), microcode read-only memories (ROMs), etc. In one embodiment, instruction cache unit 434 may be further coupled to a level 2 (L2) cache unit 476 in memory unit 470. Decode unit 440 may be coupled to a rename/allocator unit 452 in execution engine unit 450.

Execution engine unit 450 may include rename/allocator unit 452 coupled to a retirement unit 454 and a set of one or more scheduler units 456. Scheduler units 456 represent any number of different schedulers, including reservations stations, central instruction window, etc. Scheduler units 456 may be coupled to physical register file units 458. Each of physical register file units 458 represents one or more physical register files, different ones of which store one or more different data types, such as scalar integer, scalar floating point, packed integer, packed floating point, vector integer, vector floating point, etc., status (e.g., an instruction pointer that is the address of the next instruction to be executed), etc. Physical register file units 458 may be overlapped by retirement unit 154 to illustrate various ways in which register renaming and out-of-order execution may be implemented (e.g., using one or more reorder buffers and one or more retirement register files, using one or more future files, one or more history buffers, and one or more retirement register files; using register maps and a pool of registers; etc.). Generally, the architectural registers may be visible from the outside of the processor or from a programmer's perspective. The registers might not be limited to any known particular type of circuit. Various different types of registers may be suitable as long as they store and provide data as described herein. Examples of suitable registers include, but might not be limited to, dedicated physical registers, dynamically allocated physical registers using register renaming, combinations of dedicated and dynamically allocated physical registers, etc. Retirement unit 454 and physical register file units

## 16

458 may be coupled to execution clusters 460. Execution clusters 460 may include a set of one or more execution units 162 and a set of one or more memory access units 464. Execution units 462 may perform various operations (e.g., shifts, addition, subtraction, multiplication) and on various types of data (e.g., scalar floating point, packed integer, packed floating point, vector integer, vector floating point). While some embodiments may include a number of execution units dedicated to specific functions or sets of functions, other embodiments may include only one execution unit or multiple execution units that all perform all functions. Scheduler units 456, physical register file units 458, and execution clusters 460 are shown as being possibly plural because certain embodiments create separate pipelines for certain types of data/operations (e.g., a scalar integer pipeline, a scalar floating point/packed integer/packed floating point/vector integer/vector floating point pipeline, and/or a memory access pipeline that each have their own scheduler unit, physical register file unit, and/or execution cluster—and in the case of a separate memory access pipeline, certain embodiments may be implemented in which only the execution cluster of this pipeline has memory access units 464). It should also be understood that where separate pipelines are used, one or more of these pipelines may be out-of-order issue/execution and the rest in-order.

The set of memory access units 464 may be coupled to memory unit 470, which may include a data TLB unit 472 coupled to a data cache unit 474 coupled to a level 2 (L2) cache unit 476. In one exemplary embodiment, memory access units 464 may include a load unit, a store address unit, and a store data unit, each of which may be coupled to data TLB unit 472 in memory unit 470. L2 cache unit 476 may be coupled to one or more other levels of cache and eventually to a main memory.

By way of example, the exemplary register renaming, out-of-order issue/execution core architecture may implement pipeline 400 as follows: 1) instruction fetch 438 may perform fetch and length decoding stages 402 and 404; 2) decode unit 440 may perform decode stage 406; 3) rename/allocator unit 452 may perform allocation stage 408 and renaming stage 410; 4) scheduler units 456 may perform schedule stage 412; 5) physical register file units 458 and memory unit 470 may perform register read/memory read stage 414; execution cluster 460 may perform execute stage 416; 6) memory unit 470 and physical register file units 458 may perform write-back/memory-write stage 418; 7) various units may be involved in the performance of exception handling stage 422; and 8) retirement unit 454 and physical register file units 458 may perform commit stage 424.

Core 490 may support one or more instructions sets (e.g., the x86 instruction set (with some extensions that have been added with newer versions); the MIPS instruction set of MIPS Technologies of Sunnyvale, Calif.; the ARM instruction set (with optional additional extensions such as NEON) of ARM Holdings of Sunnyvale, Calif.).

It should be understood that the core may support multithreading (executing two or more parallel sets of operations or threads) in a variety of manners. Multithreading support may be performed by, for example, including time sliced multithreading, simultaneous multithreading (where a single physical core provides a logical core for each of the threads that physical core is simultaneously multithreading), or a combination thereof. Such a combination may include, for example, time sliced fetching and decoding and simultaneous multithreading thereafter such as in the Intel® Hyperthreading technology.



While register renaming may be described in the context of out-of-order execution, it should be understood that register renaming may be used in an in-order architecture. While the illustrated embodiment of the processor may also include a separate instruction and data cache units **434/474** and a shared L2 cache unit **476**, other embodiments may have a single internal cache for both instructions and data, such as, for example, a Level 1 (L1) internal cache, or multiple levels of internal cache. In some embodiments, the system may include a combination of an internal cache and an external cache that may be external to the core and/or the processor. In other embodiments, all of the cache may be external to the core and/or the processor.

FIG. 5A is a block diagram of a processor **500**, in accordance with embodiments of the present disclosure. In one embodiment, processor **500** may include a multicore processor. Processor **500** may include a system agent **510** communicatively coupled to one or more cores **502**. Furthermore, cores **502** and system agent **510** may be communicatively coupled to one or more caches **506**. Cores **502**, system agent **510**, and caches **506** may be communicatively coupled via one or more memory control units **552**. Furthermore, cores **502**, system agent **510**, and caches **506** may be communicatively coupled to a graphics module **560** via memory control units **552**.

Processor **500** may include any suitable mechanism for interconnecting cores **502**, system agent **510**, and caches **506**, and graphics module **560**. In one embodiment, processor **500** may include a ring-based interconnect unit **508** to interconnect cores **502**, system agent **510**, and caches **506**, and graphics module **560**. In other embodiments, processor **500** may include any number of well-known techniques for interconnecting such units. Ring-based interconnect unit **508** may utilize memory control units **552** to facilitate interconnections.

Processor **500** may include a memory hierarchy comprising one or more levels of caches within the cores, one or more shared cache units such as caches **506**, or external memory (not shown) coupled to the set of integrated memory controller units **552**. Caches **506** may include any suitable cache. In one embodiment, caches **506** may include one or more mid-level caches, such as level 2 (L2), level 3 (L3), level 4 (L4), or other levels of cache, a last level cache (LLC), and/or combinations thereof.

In various embodiments, one or more of cores **502** may perform multithreading. System agent **510** may include components for coordinating and operating cores **502**. System agent unit **510** may include for example a power control unit (PCU). The PCU may be or include logic and components needed for regulating the power state of cores **502**. System agent **510** may include a display engine **512** for driving one or more externally connected displays or graphics module **560**. System agent **510** may include an interface **1214** for communications busses for graphics. In one embodiment, interface **1214** may be implemented by PCI Express (PCIe). In a further embodiment, interface **1214** may be implemented by PCI Express Graphics (PEG). System agent **510** may include a direct media interface (DMI) **516**. DMI **516** may provide links between different bridges on a motherboard or other portion of a computer system. System agent **510** may include a PCIe bridge **1218** for providing PCIe links to other elements of a computing system. PCIe bridge **1218** may be implemented using a memory controller **1220** and coherence logic **1222**.

Cores **502** may be implemented in any suitable manner. Cores **502** may be homogenous or heterogeneous in terms of architecture and/or instruction set. In one embodiment, some

of cores **502** may be in-order while others may be out-of-order. In another embodiment, two or more of cores **502** may execute the same instruction set, while others may execute only a subset of that instruction set or a different instruction set.

Processor **500** may include a general-purpose processor, such as a Core™ i3, i5, i7, 2 Duo and Quad, Xeon™, Itanium™, XScale™ or StrongARM™ processor, which may be available from Intel Corporation, of Santa Clara, Calif. Processor **500** may be provided from another company, such as ARM Holdings, Ltd, MIPS, etc. Processor **500** may be a special-purpose processor, such as, for example, a network or communication processor, compression engine, graphics processor, co-processor, embedded processor, or the like. Processor **500** may be implemented on one or more chips. Processor **500** may be a part of and/or may be implemented on one or more substrates using any of a number of process technologies, such as, for example, BiCMOS, CMOS, or NMOS.

In one embodiment, a given one of caches **506** may be shared by multiple ones of cores **502**. In another embodiment, a given one of caches **506** may be dedicated to one of cores **502**. The assignment of caches **506** to cores **502** may be handled by a cache controller or other suitable mechanism. A given one of caches **506** may be shared by two or more cores **502** by implementing time-slices of a given cache **506**.

Graphics module **560** may implement an integrated graphics processing subsystem. In one embodiment, graphics module **560** may include a graphics processor. Furthermore, graphics module **560** may include a media engine **565**. Media engine **565** may provide media encoding and video decoding.

FIG. 5B is a block diagram of an example implementation of a core **502**, in accordance with embodiments of the present disclosure. Core **502** may include a front end **570** communicatively coupled to an out-of-order engine **580**. Core **502** may be communicatively coupled to other portions of processor **500** through cache hierarchy **503**.

Front end **570** may be implemented in any suitable manner, such as fully or in part by front end **201** as described above. In one embodiment, front end **570** may communicate with other portions of processor **500** through cache hierarchy **503**. In a further embodiment, front end **570** may fetch instructions from portions of processor **500** and prepare the instructions to be used later in the processor pipeline as they are passed to out-of-order execution engine **580**.

Out-of-order execution engine **580** may be implemented in any suitable manner, such as fully or in part by out-of-order execution engine **203** as described above. Out-of-order execution engine **580** may prepare instructions received from front end **570** for execution. Out-of-order execution engine **580** may include an allocate module **1282**. In one embodiment, allocate module **1282** may allocate resources of processor **500** or other resources, such as registers or buffers, to execute a given instruction. Allocate module **1282** may make allocations in schedulers, such as a memory scheduler, fast scheduler, or floating point scheduler. Such schedulers may be represented in FIG. 5B by resource schedulers **584**. Allocate module **1282** may be implemented fully or in part by the allocation logic described in conjunction with FIG. 2. Resource schedulers **584** may determine when an instruction is ready to execute based on the readiness of a given resource's sources and the availability of execution resources needed to execute an instruction. Resource schedulers **584** may be implemented by, for example, schedulers **202**, **204**, **206** as discussed above. Resource schedulers **584** may schedule the execution of instructions upon one or more resources. In one embodiment, such resources may be internal to core



**502**, and may be illustrated, for example, as resources **586**. In another embodiment, such resources may be external to core **502** and may be accessible by, for example, cache hierarchy **503**. Resources may include, for example, memory, caches, register files, or registers. Resources internal to core **502** may be represented by resources **586** in FIG. **5B**. As necessary, values written to or read from resources **586** may be coordinated with other portions of processor **500** through, for example, cache hierarchy **503**. As instructions are assigned resources, they may be placed into a reorder buffer **588**. Reorder buffer **588** may track instructions as they are executed and may selectively reorder their execution based upon any suitable criteria of processor **500**. In one embodiment, reorder buffer **588** may identify instructions or a series of instructions that may be executed independently. Such instructions or a series of instructions may be executed in parallel from other such instructions. Parallel execution in core **502** may be performed by any suitable number of separate execution blocks or virtual processors. In one embodiment, shared resources—such as memory, registers, and caches—may be accessible to multiple virtual processors within a given core **502**. In other embodiments, shared resources may be accessible to multiple processing entities within processor **500**.

Cache hierarchy **503** may be implemented in any suitable manner. For example, cache hierarchy **503** may include one or more lower or mid-level caches, such as caches **572**, **574**. In one embodiment, cache hierarchy **503** may include an LLC **595** communicatively coupled to caches **572**, **574**. In another embodiment, LLC **595** may be implemented in a module **590** accessible to all processing entities of processor **500**. In a further embodiment, module **590** may be implemented in an uncore module of processors from Intel, Inc. Module **590** may include portions or subsystems of processor **500** necessary for the execution of core **502** but might not be implemented within core **502**. Besides LLC **595**, Module **590** may include, for example, hardware interfaces, memory coherency coordinators, interprocessor interconnects, instruction pipelines, or memory controllers. Access to RAM **599** available to processor **500** may be made through module **590** and, more specifically, LLC **595**. Furthermore, other instances of core **502** may similarly access module **590**. Coordination of the instances of core **502** may be facilitated in part through module **590**.

FIGS. **6-8** may illustrate exemplary systems suitable for including processor **500**, while FIG. **9** may illustrate an exemplary system on a chip (SoC) that may include one or more of cores **502**. Other system designs and implementations known in the arts for laptops, desktops, handheld PCs, personal digital assistants, engineering workstations, servers, network devices, network hubs, switches, embedded processors, digital signal processors (DSPs), graphics devices, video game devices, set-top boxes, micro controllers, cell phones, portable media players, hand held devices, and various other electronic devices, may also be suitable. In general, a huge variety of systems or electronic devices that incorporate a processor and/or other execution logic as disclosed herein may be generally suitable.

FIG. **6** illustrates a block diagram of a system **600**, in accordance with embodiments of the present disclosure. System **600** may include one or more processors **610**, **615**, which may be coupled to graphics memory controller hub (GMCH) **620**. The optional nature of additional processors **615** is denoted in FIG. **6** with broken lines.

Each processor **610**, **615** may be some version of processor **500**. However, it should be noted that integrated graphics logic and integrated memory control units might not exist in

processors **610**, **615**. FIG. **6** illustrates that GMCH **620** may be coupled to a memory **640** that may be, for example, a dynamic random access memory (DRAM). The DRAM may, for at least one embodiment, be associated with a non-volatile cache.

GMCH **620** may be a chipset, or a portion of a chipset. GMCH **620** may communicate with processors **610**, **615** and control interaction between processors **610**, **615** and memory **640**. GMCH **620** may also act as an accelerated bus interface between the processors **610**, **615** and other elements of system **600**. In one embodiment, GMCH **620** communicates with processors **610**, **615** via a multi-drop bus, such as a frontside bus (FSB) **695**.

Furthermore, GMCH **620** may be coupled to a display **645** (such as a flat panel display). In one embodiment, GMCH **620** may include an integrated graphics accelerator. GMCH **620** may be further coupled to an input/output (I/O) controller hub (ICH) **650**, which may be used to couple various peripheral devices to system **600**. External graphics device **660** may include be a discrete graphics device coupled to ICH **650** along with another peripheral device **670**.

In other embodiments, additional or different processors may also be present in system **600**. For example, additional processors **610**, **615** may include additional processors that may be the same as processor **610**, additional processors that may be heterogeneous or asymmetric to processor **610**, accelerators (such as, e.g., graphics accelerators or digital signal processing (DSP) units), field programmable gate arrays, or any other processor. There may be a variety of differences between the physical resources **610**, **615** in terms of a spectrum of metrics of merit including architectural, micro-architectural, thermal, power consumption characteristics, and the like. These differences may effectively manifest themselves as asymmetry and heterogeneity amongst processors **610**, **615**. For at least one embodiment, various processors **610**, **615** may reside in the same die package.

FIG. **7** illustrates a block diagram of a second system **700**, in accordance with embodiments of the present disclosure. As shown in FIG. **7**, multiprocessor system **700** may include a point-to-point interconnect system, and may include a first processor **770** and a second processor **780** coupled via a point-to-point interconnect **750**. Each of processors **770** and **780** may be some version of processor **500** as one or more of processors **610**, **615**.

While FIG. **7** may illustrate two processors **770**, **780**, it is to be understood that the scope of the present disclosure is not so limited. In other embodiments, one or more additional processors may be present in a given processor.

Processors **770** and **780** are shown including integrated memory controller units **772** and **782**, respectively. Processor **770** may also include as part of its bus controller units point-to-point (P-P) interfaces **776** and **778**; similarly, second processor **780** may include P-P interfaces **786** and **788**. Processors **770**, **780** may exchange information via a point-to-point (P-P) interface **750** using P-P interface circuits **778**, **788**. As shown in FIG. **7**, IMCs **772** and **782** may couple the processors to respective memories, namely a memory **732** and a memory **734**, which in one embodiment may be portions of main memory locally attached to the respective processors.

Processors **770**, **780** may each exchange information with a chipset **790** via individual P-P interfaces **752**, **754** using point to point interface circuits **776**, **794**, **786**, **798**. In one embodiment, chipset **790** may also exchange information with a high-performance graphics circuit **738** via a high-performance graphics interface **739**.

A shared cache (not shown) may be included in either processor or outside of both processors, yet connected with

## 21

the processors via P-P interconnect, such that either or both processors' local cache information may be stored in the shared cache if a processor is placed into a low power mode.

Chipset **790** may be coupled to a first bus **716** via an interface **796**. In one embodiment, first bus **716** may be a Peripheral Component Interconnect (PCI) bus, or a bus such as a PCI Express bus or another third generation I/O interconnect bus, although the scope of the present disclosure is not so limited.

As shown in FIG. 7, various I/O devices **714** may be coupled to first bus **716**, along with a bus bridge **718** which couples first bus **716** to a second bus **720**. In one embodiment, second bus **720** may be a low pin count (LPC) bus. Various devices may be coupled to second bus **720** including, for example, a keyboard and/or mouse **722**, communication devices **727** and a storage unit **728** such as a disk drive or other mass storage device which may include instructions/code and data **730**, in one embodiment. Further, an audio I/O **724** may be coupled to second bus **720**. Note that other architectures may be possible. For example, instead of the point-to-point architecture of FIG. 7, a system may implement a multi-drop bus or other such architecture.

FIG. 8 illustrates a block diagram of a third system **800** in accordance with embodiments of the present disclosure. Like elements in FIGS. 7 and 8 bear like reference numerals, and certain aspects of FIG. 7 have been omitted from FIG. 8 in order to avoid obscuring other aspects of FIG. 8.

FIG. 8 illustrates that processors **870**, **880** may include integrated memory and I/O control logic ("CL") **872** and **882**, respectively. For at least one embodiment, CL **872**, **882** may include integrated memory controller units such as that described above in connection with FIGS. 5 and 7. In addition, CL **872**, **882** may also include I/O control logic. FIG. 8 illustrates that not only memories **832**, **834** may be coupled to CL **872**, **882**, but also that I/O devices **814** may also be coupled to control logic **872**, **882**. Legacy I/O devices **815** may be coupled to chipset **890**.

FIG. 9 illustrates a block diagram of a SoC **900**, in accordance with embodiments of the present disclosure. Similar elements in FIG. 5 bear like reference numerals. Also, dashed lined boxes may represent optional features on more advanced SoCs. An interconnect units **902** may be coupled to: an application processor **910** which may include a set of one or more cores **902A-N** and shared cache units **906**; a system agent unit **910**; a bus controller units **916**; an integrated memory controller units **914**; a set of one or more media processors **920** which may include integrated graphics logic **908**, an image processor **924** for providing still and/or video camera functionality, an audio processor **926** for providing hardware audio acceleration, and a video processor **928** for providing video encode/decode acceleration; an static random access memory (SRAM) unit **930**; a direct memory access (DMA) unit **932**; and a display unit **940** for coupling to one or more external displays.

FIG. 10 illustrates a processor containing a central processing unit (CPU) and a graphics processing unit (GPU), which may perform at least one instruction, in accordance with embodiments of the present disclosure. In one embodiment, an instruction to perform operations according to at least one embodiment could be performed by the CPU. In another embodiment, the instruction could be performed by the GPU. In still another embodiment, the instruction may be performed through a combination of operations performed by the GPU and the CPU. For example, in one embodiment, an instruction in accordance with one embodiment may be received and decoded for execution on the GPU. However, one or more operations within the decoded instruction may be

## 22

performed by a CPU and the result returned to the GPU for final retirement of the instruction. Conversely, in some embodiments, the CPU may act as the primary processor and the GPU as the co-processor.

In some embodiments, instructions that benefit from highly parallel, throughput processors may be performed by the GPU, while instructions that benefit from the performance of processors that benefit from deeply pipelined architectures may be performed by the CPU. For example, graphics, scientific applications, financial applications and other parallel workloads may benefit from the performance of the GPU and be executed accordingly, whereas more sequential applications, such as operating system kernel or application code may be better suited for the CPU.

In FIG. 10, processor **1000** includes a CPU **1005**, GPU **1010**, image processor **1015**, video processor **1020**, USB controller **1025**, UART controller **1030**, SPI/SDIO controller **1035**, display device **1040**, memory interface controller **1045**, MIPI controller **1050**, flash memory controller **1055**, dual data rate (DDR) controller **1060**, security engine **1065**, and I<sup>2</sup>S/I<sup>2</sup>C controller **1070**. Other logic and circuits may be included in the processor of FIG. 10, including more CPUs or GPUs and other peripheral interface controllers.

One or more aspects of at least one embodiment may be implemented by representative data stored on a machine-readable medium which represents various logic within the processor, which when read by a machine causes the machine to fabricate logic to perform the techniques described herein. Such representations, known as "IP cores" may be stored on a tangible, machine-readable medium ("tape") and supplied to various customers or manufacturing facilities to load into the fabrication machines that actually make the logic or processor. For example, IP cores, such as the Cortex™ family of processors developed by ARM Holdings, Ltd. and Loongson IP cores developed the Institute of Computing Technology (ICT) of the Chinese Academy of Sciences may be licensed or sold to various customers or licensees, such as Texas Instruments, Qualcomm, Apple, or Samsung and implemented in processors produced by these customers or licensees.

FIG. 11 illustrates a block diagram illustrating the development of IP cores, in accordance with embodiments of the present disclosure. Storage **1130** may include simulation software **1120** and/or hardware or software model **1110**. In one embodiment, the data representing the IP core design may be provided to storage **1130** via memory **1140** (e.g., hard disk), wired connection (e.g., internet) **1150** or wireless connection **1160**. The IP core information generated by the simulation tool and model may then be transmitted to a fabrication facility where it may be fabricated by a 3<sup>rd</sup> party to perform at least one instruction in accordance with at least one embodiment.

In some embodiments, one or more instructions may correspond to a first type or architecture (e.g., x86) and be translated or emulated on a processor of a different type or architecture (e.g., ARM). An instruction, according to one embodiment, may therefore be performed on any processor or processor type, including ARM, x86, MIPS, a GPU, or other processor type or architecture.

FIG. 12 illustrates how an instruction of a first type may be emulated by a processor of a different type, in accordance with embodiments of the present disclosure. In FIG. 12, program **1205** contains some instructions that may perform the same or substantially the same function as an instruction according to one embodiment. However the instructions of program **1205** may be of a type and/or format that is different from or incompatible with processor **1215**, meaning the instructions of the type in program **1205** may not be able to

23

execute natively by the processor **1215**. However, with the help of emulation logic, **1210**, the instructions of program **1205** may be translated into instructions that may be natively executed by the processor **1215**. In one embodiment, the emulation logic may be embodied in hardware. In another embodiment, the emulation logic may be embodied in a tangible, machine-readable medium containing software to translate instructions of the type in program **1205** into the type natively executable by processor **1215**. In other embodiments, emulation logic may be a combination of fixed-function or programmable hardware and a program stored on a tangible, machine-readable medium. In one embodiment, the processor contains the emulation logic, whereas in other embodiments, the emulation logic exists outside of the processor and may be provided by a third party. In one embodiment, the processor may load the emulation logic embodied in a tangible, machine-readable medium containing software by executing microcode or firmware contained in or associated with the processor.

FIG. **13** is a block diagram contrasting the use of a software instruction converter to convert binary instructions in a source instruction set to binary instructions in a target instruction set according to embodiments of the invention. In the illustrated embodiment, the instruction converter is a software instruction converter, although alternatively the instruction converter may be implemented in software, firmware, hardware, or various combinations thereof. FIG. **13** shows a program in a high level language **1302** may be compiled using an x86 compiler **1304** to generate x86 binary code **1306** that may be natively executed by a processor with at least one x86 instruction set core **1316**. The processor with at least one x86 instruction set core **1316** represents any processor that can perform substantially the same functions as an Intel processor with at least one x86 instruction set core by compatibly executing or otherwise processing (1) a substantial portion of the instruction set of the Intel x86 instruction set core or (2) object code versions of applications or other software targeted to run on an Intel processor with at least one x86 instruction set core, in order to achieve substantially the same result as an Intel processor with at least one x86 instruction set core. The x86 compiler **1304** represents a compiler that is operable to generate x86 binary code **1306** (e.g., object code) that can, with or without additional linkage processing, be executed on the processor with at least one x86 instruction set core **1316**. Similarly, FIG. **13** shows the program in the high level language **1302** may be compiled using an alternative instruction set compiler **1308** to generate alternative instruction set binary code **1310** that may be natively executed by a processor without at least one x86 instruction set core **1314** (e.g., a processor with cores that execute the MIPS instruction set of MIPS Technologies of Sunnyvale, Calif. and/or that execute the ARM instruction set of ARM Holdings of Sunnyvale, Calif.).

The instruction converter **1312** is used to convert the x86 binary code **1306** into alternative instruction set binary code **1311** that may be natively executed by the processor without an x86 instruction set core **1314**. This converted code may or may not be the same as the alternative instruction set binary code **1310** resulting from an alternative instruction set compiler **1308**; however, the converted code will accomplish the same general operation and be made up of instructions from the alternative instruction set. Thus, the instruction converter **1312** represents software, firmware, hardware, or a combination thereof that, through emulation, simulation or any other process, allows a processor or other electronic device that does not have an x86 instruction set processor or core to execute the x86 binary code **1306**.

24

FIG. **14** is a block diagram of an instruction set architecture **1400** of a processor, in accordance with embodiments of the present disclosure. Instruction set architecture **1400** may include any suitable number or kind of components.

For example, instruction set architecture **1400** may include processing entities such as one or more cores **1406**, **1407** and a graphics processing unit **1415**. Cores **1406**, **1407** may be communicatively coupled to the rest of instruction set architecture **1400** through any suitable mechanism, such as through a bus or cache. In one embodiment, cores **1406**, **1407** may be communicatively coupled through an L2 cache control **1408**, which may include a bus interface unit **1409** and an L2 cache **1410**. Cores **1406**, **1407** and graphics processing unit **1415** may be communicatively coupled to each other and to the remainder of instruction set architecture **1400** through interconnect **1410**. In one embodiment, graphics processing unit **1415** may use a video code **1420** defining the manner in which particular video signals will be encoded and decoded for output.

Instruction set architecture **1400** may also include any number or kind of interfaces, controllers, or other mechanisms for interfacing or communicating with other portions of an electronic device or system. Such mechanisms may facilitate interaction with, for example, peripherals, communications devices, other processors, or memory. In the example of FIG. **14**, instruction set architecture **1400** may include a liquid crystal display (LCD) video interface **1425**, a subscriber interface module (SIM) interface **1430**, a boot ROM interface **1435**, a synchronous dynamic random access memory (SDRAM) controller **1440**, a flash controller **1445**, and a serial peripheral interface (SPI) master unit **1450**. LCD video interface **1425** may provide output of video signals from, for example, GPU **1415** and through, for example, a mobile industry processor interface (MIPI) **1490** or a high-definition multimedia interface (HDMI) **1495** to a display. Such a display may include, for example, an LCD. SIM interface **1430** may provide access to or from a SIM card or device. SDRAM controller **1440** may provide access to or from memory such as an SDRAM chip or module. Flash controller **1445** may provide access to or from memory such as flash memory or other instances of RAM. SPI master unit **1450** may provide access to or from communications modules, such as a Bluetooth module **1470**, high-speed 3G modem **1475**, global positioning system module **1480**, or wireless module **1485** implementing a communications standard such as 802.11.

FIG. **15** is a more detailed block diagram of an instruction set architecture **1500** of a processor, in accordance with embodiments of the present disclosure. Instruction architecture **1500** may implement one or more aspects of instruction set architecture **1400**. Furthermore, instruction set architecture **1500** may illustrate modules and mechanisms for the execution of instructions within a processor.

Instruction architecture **1500** may include a memory system **1540** communicatively coupled to one or more execution entities **1565**. Furthermore, instruction architecture **1500** may include a caching and bus interface unit such as unit **1510** communicatively coupled to execution entities **1565** and memory system **1540**. In one embodiment, loading of instructions into execution entities **1564** may be performed by one or more stages of execution. Such stages may include, for example, instruction prefetch stage **1530**, dual instruction decode stage **1550**, register rename stage **155**, issue stage **1560**, and writeback stage **1570**.

In another embodiment, memory system **1540** may include a retirement pointer **1582**. Retirement pointer **1582** may store a value identifying the program order (PO) of the last retired

## 25

instruction. Retirement pointer **1582** may be set by, for example, retirement unit **454**. If no instructions have yet been retired, retirement pointer **1582** may include a null value.

Execution entities **1565** may include any suitable number and kind of mechanisms by which a processor may execute instructions. In the example of FIG. **15**, execution entities **1565** may include ALU/multiplication units (MUL.) **1566**, ALUs **1567**, and floating point units (FPU) **1568**. In one embodiment, such entities may make use of information contained within a given address **1569**. Execution entities **1565** in combination with stages **1530**, **1550**, **1555**, **1560**, **1570** may collectively form an execution unit.

Unit **1510** may be implemented in any suitable manner. In one embodiment, unit **1510** may perform cache control. In such an embodiment, unit **1510** may thus include a cache **1525**. Cache **1525** may be implemented, in a further embodiment, as an L2 unified cache with any suitable size, such as zero, 128k, 256k, 512k, 1M, or 2M bytes of memory. In another, further embodiment, cache **1525** may be implemented in error-correcting code memory. In another embodiment, unit **1510** may perform bus interfacing to other portions of a processor or electronic device. In such an embodiment, unit **1510** may thus include a bus interface unit **1520** for communicating over an interconnect, intraprocessor bus, interprocessor bus, or other communication bus, port, or line. Bus interface unit **1520** may provide interfacing in order to perform, for example, generation of the memory and input/output addresses for the transfer of data between execution entities **1565** and the portions of a system external to instruction architecture **1500**.

To further facilitate its functions, bus interface unit **1520** may include an interrupt control and distribution unit **1511** for generating interrupts and other communications to other portions of a processor or electronic device. In one embodiment, bus interface unit **1520** may include a snoop control unit **1512** that handles cache access and coherency for multiple processing cores. In a further embodiment, to provide such functionality, snoop control unit **1512** may include a cache-to-cache transfer unit that handles information exchanges between different caches. In another, further embodiment, snoop control unit **1512** may include one or more snoop filters **1514** that monitors the coherency of other caches (not shown) so that a cache controller, such as unit **1510**, does not have to perform such monitoring directly. Unit **1510** may include any suitable number of timers **1515** for synchronizing the actions of instruction architecture **1500**. Also, unit **1510** may include an AC port **1516**.

Memory system **1540** may include any suitable number and kind of mechanisms for storing information for the processing needs of instruction architecture **1500**. In one embodiment, memory system **1504** may include a load store unit **1530** for storing information such as buffers written to or read back from memory or registers. In another embodiment, memory system **1504** may include a translation lookaside buffer (TLB) **1545** that provides look-up of address values between physical and virtual addresses. In yet another embodiment, bus interface unit **1520** may include a memory management unit (MMU) **1544** for facilitating access to virtual memory. In still yet another embodiment, memory system **1504** may include a prefetcher **1543** for requesting instructions from memory before such instructions are actually needed to be executed, in order to reduce latency.

The operation of instruction architecture **1500** to execute an instruction may be performed through different stages. For example, using unit **1510** instruction prefetch stage **1530** may access an instruction through prefetcher **1543**. Instructions retrieved may be stored in instruction cache **1532**. Prefetch

## 26

stage **1530** may enable an option **1531** for fast-loop mode, wherein a series of instructions forming a loop that is small enough to fit within a given cache are executed. In one embodiment, such an execution may be performed without needing to access additional instructions from, for example, instruction cache **1532**. Determination of what instructions to prefetch may be made by, for example, branch prediction unit **1535**, which may access indications of execution in global history **1536**, indications of target addresses **1537**, or contents of a return stack **1538** to determine which of branches **1557** of code will be executed next. Such branches may be possibly prefetched as a result. Branches **1557** may be produced through other stages of operation as described below. Instruction prefetch stage **1530** may provide instructions as well as any predictions about future instructions to dual instruction decode stage.

Dual instruction decode stage **1550** may translate a received instruction into microcode-based instructions that may be executed. Dual instruction decode stage **1550** may simultaneously decode two instructions per clock cycle. Furthermore, dual instruction decode stage **1550** may pass its results to register rename stage **1555**. In addition, dual instruction decode stage **1550** may determine any resulting branches from its decoding and eventual execution of the microcode. Such results may be input into branches **1557**.

Register rename stage **1555** may translate references to virtual registers or other resources into references to physical registers or resources. Register rename stage **1555** may include indications of such mapping in a register pool **1556**. Register rename stage **1555** may alter the instructions as received and send the result to issue stage **1560**.

Issue stage **1560** may issue or dispatch commands to execution entities **1565**. Such issuance may be performed in an out-of-order fashion. In one embodiment, multiple instructions may be held at issue stage **1560** before being executed. Issue stage **1560** may include an instruction queue **1561** for holding such multiple commands. Instructions may be issued by issue stage **1560** to a particular processing entity **1565** based upon any acceptable criteria, such as availability or suitability of resources for execution of a given instruction. In one embodiment, issue stage **1560** may reorder the instructions within instruction queue **1561** such that the first instructions received might not be the first instructions executed. Based upon the ordering of instruction queue **1561**, additional branching information may be provided to branches **1557**. Issue stage **1560** may pass instructions to executing entities **1565** for execution.

Upon execution, writeback stage **1570** may write data into registers, queues, or other structures of instruction set architecture **1500** to communicate the completion of a given command. Depending upon the order of instructions arranged in issue stage **1560**, the operation of writeback stage **1570** may enable additional instructions to be executed. Performance of instruction set architecture **1500** may be monitored or debugged by trace unit **1575**.

FIG. **16** is a block diagram of an execution pipeline **1600** for an instruction set architecture of a processor, in accordance with embodiments of the present disclosure. Execution pipeline **1600** may illustrate operation of, for example, instruction architecture **1500** of FIG. **15**.

Execution pipeline **1600** may include any suitable combination of steps or operations. In **1605**, predictions of the branch that is to be executed next may be made. In one embodiment, such predictions may be based upon previous executions of instructions and the results thereof. In **1610**, instructions corresponding to the predicted branch of execution may be loaded into an instruction cache. In **1615**, one or

more such instructions in the instruction cache may be fetched for execution. In **1620**, the instructions that have been fetched may be decoded into microcode or more specific machine language. In one embodiment, multiple instructions may be simultaneously decoded. In **1625**, references to registers or other resources within the decoded instructions may be reassigned. For example, references to virtual registers may be replaced with references to corresponding physical registers. In **1630**, the instructions may be dispatched to queues for execution. In **1640**, the instructions may be executed. Such execution may be performed in any suitable manner. In **1650**, the instructions may be issued to a suitable execution entity. The manner in which the instruction is executed may depend upon the specific entity executing the instruction. For example, at **1655**, an ALU may perform arithmetic functions. The ALU may utilize a single clock cycle for its operation, as well as two shifters. In one embodiment, two ALUs may be employed, and thus two instructions may be executed at **1655**. At **1660**, a determination of a resulting branch may be made. A program counter may be used to designate the destination to which the branch will be made. **1660** may be executed within a single clock cycle. At **1665**, floating point arithmetic may be performed by one or more FPUs. The floating point operation may require multiple clock cycles to execute, such as two to ten cycles. At **1670**, multiplication and division operations may be performed. Such operations may be performed in four clock cycles. At **1675**, loading and storing operations to registers or other portions of pipeline **1600** may be performed. The operations may include loading and storing addresses. Such operations may be performed in four clock cycles. At **1680**, write-back operations may be performed as required by the resulting operations of **1655-1675**.

FIG. **17** is a block diagram of an electronic device **1700** for utilizing a processor **1710**, in accordance with embodiments of the present disclosure. Electronic device **1700** may include, for example, a notebook, an ultrabook, a computer, a tower server, a rack server, a blade server, a laptop, a desktop, a tablet, a mobile device, a phone, an embedded computer, or any other suitable electronic device.

Electronic device **1700** may include processor **1710** communicatively coupled to any suitable number or kind of components, peripherals, modules, or devices. Such coupling may be accomplished by any suitable kind of bus or interface, such as I<sup>2</sup>C bus, system management bus (SMBus), low pin count (LPC) bus, SPI, high definition audio (HDA) bus, Serial Advance Technology Attachment (SATA) bus, USB bus (versions 1, 2, 3), or Universal Asynchronous Receiver/Transmitter (UART) bus.

Such components may include, for example, a display **1724**, a touch screen **1725**, a touch pad **1730**, a near field communications (NFC) unit **1745**, a sensor hub **1740**, a thermal sensor **1746**, an express chipset (EC) **1735**, a trusted platform module (TPM) **1738**, BIOS/firmware/flash memory **1722**, a digital signal processor **1760**, a drive **1720** such as a solid state disk (SSD) or a hard disk drive (HDD), a wireless local area network (WLAN) unit **1750**, a Bluetooth unit **1752**, a wireless wide area network (WWAN) unit **1756**, a global positioning system (GPS), a camera **1754** such as a USB 3.0 camera, or a low power double data rate (LPDDR) memory unit **1715** implemented in, for example, the LPDDR3 standard. These components may each be implemented in any suitable manner.

Furthermore, in various embodiments other components may be communicatively coupled to processor **1710** through the components discussed above. For example, an accelerometer **1741**, ambient light sensor (ALS) **1742**, compass **1743**,

and gyroscope **1744** may be communicatively coupled to sensor hub **1740**. A thermal sensor **1739**, fan **1737**, keyboard **1746**, and touch pad **1730** may be communicatively coupled to EC **1735**. Speaker **1763**, headphones **1764**, and a microphone **1765** may be communicatively coupled to an audio unit **1764**, which may in turn be communicatively coupled to DSP **1760**. Audio unit **1764** may include, for example, an audio codec and a class D amplifier. A SIM card **1757** may be communicatively coupled to WWAN unit **1756**. Components such as WLAN unit **1750** and Bluetooth unit **1752**, as well as WWAN unit **1756** may be implemented in a next generation form factor (NGFF).

Embodiments of the present disclosure involve an instruction and logic for adaptive dataset priorities. Such priorities may be used to evaluate and evict elements from caches. FIG. **18** is a block diagram of a system **1800** for implementing an instruction and logic for adaptive dataset priorities, in accordance with embodiments of the present disclosure. In one embodiment, dataset priorities may be assigned to memory locations by a producer of such dataset priorities. The dataset priorities may be assigned to, for example, collections of objects and variables. In another embodiment and after a miss of a data access, hardware may utilize such dataset priorities to evaluate whether and how to evict data, such as cachelines, that otherwise are candidates for eviction. In a further embodiment, hardware may also consider access histories to evaluate whether and which cachelines to evict. Candidates for eviction may hold data from datasets of a high priority and the hardware may find a different candidate for eviction. In yet another embodiment, the hardware may override or adapt the dataset priorities based upon one or more attempted evictions of cachelines holding higher priority data. In still another embodiment, the hardware may collect and make available to itself and to software certain metrics of cacheline evictions to determine that software is experiencing cache-thrash. In another embodiment, the producer of the dataset priorities may take corrective action by, for example, adjusting the utilization of dataset priorities to be performed by hardware, adjusting the assigned dataset priorities, or other suitable action. The producer of dataset priorities may make such corrective action based upon a determination that there is cache thrash occurring. Furthermore, the producer of data set priorities may also initiate such corrective action based upon a determination that software is undergoing a transient or durable transition of phase.

System **1800** may include any suitable number and kind of elements to perform the operations described herein. Furthermore, although specific elements of system **1800** may be described herein as performing a specific function, any suitable portion of system **1800** may perform the functionality described herein. System **1800** may fetch, dispatch, execute, and retire instructions out-of-order.

The producer of dataset priorities may include any suitable entity to specify priorities of memory locations. Moreover, the producer of dataset priorities may employ a variety of alternative forms of identifying the datasets for purposes of associating dataset priorities. These may include, for example, specifications of addresses in logical or physical memory, in locations on disks, as objects in local or remote file systems, as keys or table locations in databases, as web addressed entities, and other reasonable means of naming entities capable of being located through software-accessible naming systems. In one embodiment, the producer of dataset priorities may be implemented in software. In another embodiment, the producer of dataset priorities may include instructions in software applications specifying the priorities. Such applications may include, for example, applications

**1810.** Applications **1810** may specify dataset priorities in terms of virtual memory, physical memory, or object identifiers in local and remote object systems. Moreover, applications **1810** may change dataset priorities, instruct hardware to perform priority analysis in a given mode, or otherwise take corrective action based upon hardware adapting to dataset priorities. In yet another embodiment, the production of dataset priorities may include instructions in operating systems. Operating systems, such as operating system **1808**, may specify dataset priorities or translate dataset priorities from virtual memory (or other object) designations to physical memory designations. Moreover, operating system **1808** may change dataset priorities, instruct hardware to perform priority analysis in a given mode, or otherwise take corrective action based upon hardware adapting to dataset priorities. In another embodiment, the production of dataset priorities may include instructions in a compiler, translator, just-in-time component, or other suitable entities in a processor **1804**. Such an entity may include a dynamic binary translator (DBT) **1816**. DBT **1816** may specify dataset priorities or translate dataset priorities from virtual memory designations to physical memory designations. Moreover, DBT **1816** may change dataset priorities, instruct hardware to perform priority analysis in a given mode, or otherwise take corrective action based upon hardware adapting to dataset priorities.

Any suitable hardware may be used to evaluate and adapt dataset priorities in view of a need to evict data. In one embodiment, such hardware may include any cache controller **1844** of processor **1804**. Furthermore, the data may need to be evicted from any suitable data container, such as a memory, cache, or buffer. The data may need to be evicted because a data access is made of the container, but the specific data requested is unavailable and the container is full. Accordingly, data in the container must be evicted to make room for the data that is to be used. For example, execution of an execution unit **1822** in a core **1820** may make a write or read of a memory location through a cache hierarchy implemented in any suitable manner. In another example, the request may be made of embedded DRAMs that are managed as transparent caches. In yet another example, a cache may reside on a peripheral interface controller such as a PCI-express storage or network adapter card. In the example of FIG. **18**, the request may be made of an L1 cache **1824**. The requested resource might not be available in L1 cache **1824**, and an associated L1 cache controller **1826** may make request the data from a higher level cache, such as L2 cache **1828**. If the data is unavailable in L2 cache **1828**, an associated L2 cache controller **1830** may make the request of an L3 cache **1832**. If the requested resource is unavailable in L3 cache **1832**, an associated cache controller **1834** may make the request directly to the associated resource, such as on-board memory **1836** or external memory **1840**. The requested data, when found in the cache hierarchy, may be propagated down through the hierarchy. Each cache hierarchy level might require the ability to evict data lines to make room for such requested data. In the example of FIG. **18**, cache controllers **1844** may represent possible operation of any suitable hardware to evaluate and adapt dataset priorities in view of a need to evict data. Such hardware may include any suitable one of cache controllers **1826**, **1830**, **1834**, or cache controllers for peripherals such as PCI-E devices like network interfaces.

System **1800** may include processor **1804** to perform, fully or in part, adaptive handling of dataset priorities in applying an eviction policy. Processor **1804** may be implemented in part by any processor core, logical processor, processor, or other processing entity such as those illustrated in FIGS. **1-17**. In various embodiments, processor **1804** may include a

front end **1812** to fetch instructions to be executed; a scheduler and allocator **1818** to allocate assign instructions for execution to execution units **1822** or cores **1820**; and one or more execution units **1822** or cores **1822** to execute the instructions. Processor **1804** may include other suitable components that are not shown, such as allocation units to reserve alias resources or retirement units to recover resources used by the instructions.

Front end **1812** may fetch and prepare instructions to be used by other elements of processor **1804**, and may include any suitable number or kind of components. For example, front end **1812** may include a decoder **1814** to translate instructions into microcode commands. Furthermore, front end **1812** may arrange instructions into parallel groups or other mechanisms of out-of-order processing. DBT **1816** may be included in front end **1812**. Instructions may be processed for in-order or out-of-order execution. In one embodiment, such instructions may include instructions for manipulating the adaptive handling of dataset priorities. In another embodiment, such instructions may include instructions for specifying dataset priorities. Scheduler **1820** may schedule instructions to be executed on any suitable execution unit **1822** or core **1822**. Cores **1822** may be implemented in any suitable manner. A given core **1822** may include any suitable number, kind, and combination of execution units **1822**.

In one embodiment, cache controller **1844** may apply eviction policies that are adaptive to information, such as priority of datasets, specified by software and to actual behavior experience in hardware. Some portions of a program's address space may be disproportionately important to its performance, efficiency, or responsiveness. In some embodiments, such importance may be recognized by software itself, either through monitoring or programming. Accordingly, software such as applications **1810**, operating system **1808**, or DBT **1816** may specify priorities as reflected in priority datasets **1806**.

In one embodiment, priority datasets **1806** may specify virtual memory address ranges and, for each range, a priority designation such as an integer value wherein a higher value indicates a higher priority. In another embodiment, priority datasets **1806** may specify physical memory address ranges and, for each range, a priority designation. In yet another embodiment, priority datasets may specify other identifiers by which data is named and located by software, before or after being brought into a computer's physical memory range. In various embodiments, applications **1810**, operating system **1808**, or DBT **1816** may generate a first priority dataset and another of applications **1810**, operating system **1808**, or DBT **1816** may translate the priority dataset into a different type. For example, application **1810** may generate a virtual memory priority dataset that is translated by operating system **1808** into a physical memory priority dataset. Hardware of processor **1804** may access priority values specified in physical memory ranges.

Some cache eviction policies may identify a candidate for eviction according to criteria such as least recently used (LRU), least frequently used (LFU), First-in First-out (FIFO), First-in-not-used-first-out (FINUFO), approximately least recently used (ALRU), or combinations thereof.

However, the use of such policies might not reliably keep the most critical data in lower level caches close to cores **1820** for faster access. In addition, full management of caches by software specifying priorities such as applications **1810**, operating system **1808**, or DBT **1816** may be too slow to effectively implement. Furthermore, such management may

31

be impossible to be fully drafted into instructions for a specific application **1810** that can consider the dynamic status of a cache.

In one embodiment, cache controllers **1844** may apply and adapt the priorities specified by software by utilizing hysteresis in victim selection. Cache controllers may allow higher priority cachelines to be less readily evicted upon misses of requests for lower priority addresses. Furthermore, cache controllers may allow a lower priority cacheline to stay in the cache only as long as the cacheline has been robustly accessed or otherwise qualified under cache eviction policies. However, software may define some address ranges with high priorities that are not used efficiently. Consequently, cache controllers **1844** may adapt and override such priorities as appropriate. The net effect may be equivalent to the software never describing the priorities initially. Furthermore, if software applications exit, and thus do not need to maintain the priority of the ranges, cache controllers **1844** may behave as if the software never provided the priorities.

Any suitable mechanism may be used to provide dataset priorities to hardware such as cache controllers **1844**. The mechanism by which dataset priorities are accessed by hardware might be required to be sufficiently fast so that performance of processor **1804** is not degraded. The hardware might be required to access the dataset priorities as a preliminary step in the cache eviction process and, accordingly, delays in such access may cause delay in cache miss handling. In one embodiment, dataset priorities may be communicated to hardware through a signature. In another embodiment, dataset priorities may be communicated to hardware through a bloom filter. Software may set priority datasets **1806** into a bloom filter **1842** at any suitable time. Moreover, software may reset priority datasets **1806** upon a change in applications, a change in contents of memory, a context switch of operating system **1808**, an attempt to reprioritize memory ranges upon feedback from hardware, or any other suitable time.

Upon determining that a cache miss or other similar request has been made, cache controllers **1844** may read the priority data from priority datasets **1806** as embodied in bloom filter **1842**. Such data, as resident within bloom filter **1842** or any other suitable structure, may yield a priority when queried with an address. The set of addresses  $A$  at a given priority  $\pi$  may be given as

$$S(\pi): \{A_i | P(A_i) = \pi\}$$

FIG. **19** is an illustration of operation of system **1800** to perform evaluation and adaptation of priority datasets during a cache miss, in accordance with embodiments of the present disclosure. In one embodiment, at (1), software may set the datasets. Such a dataset may be a mapping of priorities to ranges of virtual memory and may be performed by application **1810**, though any suitable entity may set the datasets. At (2), software may translate the dataset into a physical memory mapping, if necessary. Such translation may be performed, for example, by operating system **1808**. Moreover, the priority dataset may be entered into a suitable structure for access by the hardware, such as bloom filter **1842**. The operations of (1) and (2) may be repeated as many times as necessary and upon any suitable condition.

At (3), a cache miss or similar data request may be detected by cache controller **1844**. If the cache is full, a cache victim to be evicted might be identified. Cache controller **1844** may employ any suitable algorithm to identify a first candidate as a cache victim. Such a selection may then be evaluated or adapted according to priority datasets specified by software and upon previous searching for cache victims. Cache con-

32

troller **1844** may identify a candidate cache victim at an address  $V$ , wherein the cache victim is a cacheline with data.

In one embodiment, at (4), cache controller **1844** may identify the priority of the candidate cache victim as defined by software. Cache controller **1844** may also identify the priority of the data that caused the cache miss. Moreover, cache controller **1844** may identify the priorities by accessing bloom filter **1842**. The priority of the candidate cache victim may be compared against a threshold or against the priority of the data that caused the cache miss.

If the candidate cache victim is of a low priority, then the cache victim may be evicted. In one embodiment, at (5) if the candidate cache victim is of a high priority, then another candidate victim may be identified using the operations described in (3). Cache controller **1844** may search for a candidate with a lower priority, if one is available. At (6), the attempted access and eviction may be stored in a record **1902** such that repeated attempts to find a candidate cache victim may be evaluated. In one embodiment, cache controller **1844** may adapt its cache eviction scheme based upon such data. The effort of cache controller **1844** to find lower-priority candidates for cache eviction is thus bounded. In one embodiment, cache controller **1844** may bound the search for lower-priority candidates based upon the priority of the first or additional candidate cache victims and upon the priority of the address that is missed in the cache miss. Cache controller **1844** may thus determine whether the cache is overstuffing with high priority cachelines that are not being adequately (according to the cache victim algorithm) used in the cache. Any suitable manner of monitoring usage of high priority cachelines may be used by cache controller **1844**. In one embodiment, cache controller **1844** may utilize a moving window average of the number of failed attempts to evict high-priority cachelines. Information stored by cache controller **1844** in record **1902** may include any suitable information. In one embodiment, such information may include statistics on how much work has been performed or is needed to be performed to find low priority cachelines to evict. Such information may be shared with software so that software may recognize inadequate priority datasets and readjust them.

At (7), in one embodiment, cache controller **1844** may adapt its cache eviction scheme based upon previous attempts to find a suitable candidate cache victim. The recent history of cache victim attempts or instructions from software may indicate that the software is going through a transient phase, such as garbage collection. The priority datasets provided by software may be overridden. Furthermore, the producer of the priority datasets may be informed. At (8), software may change priorities or otherwise take corrective action at any suitable time. At (9), the identified victim may be evicted, even though the identified victim is of a higher priority identified by software.

In some embodiments, instructions for particular operations may be given negative priority to encourage their eviction. The instructions may include, for example, certain non-allocating load operations, push-store operations, or other load or store operations that carry a negative priority hint for the data that they access. In other embodiments, priority datasets may include address range of multiple applications and thus multiple address spaces.

FIG. **20** is an illustration of operation of system **1800** to specify dataset priorities, according to embodiments of the present disclosure. First, one or more address spaces **2002**, **2004** may be considered when defining dataset priorities. Each address space may belong to the same or different applications or entities. In one embodiment, each address



space may be defined according to virtual memory regions. In other embodiments address spaces may be defined recursively as enumerated regions from other address spaces. Dataset priorities for each region may be defined according to any suitable prioritization scheme or consideration of the importance of the underlying data. Such importance may be evaluated according to execution efficiency.

For example, an application may allocate an address space **2002** including six regions. Region **3** may include a texture cache and be assigned a priority of “2”. Region **6** may include a MiniMDKernel and be assigned an even higher priority of “4”. The other regions may be assigned a default priority, a priority of “0”, or remain unassigned. Hardware may treat unassigned regions as having a priority of “0”. Any numbers, positive, negative, or 0, may be used in specifying dataset priorities.

In another example, the same or a different application may allocate an address space **2004** including three regions. Region **8** may include a dictionary or other software reference and be assigned a priority of “4”. Region **10** may include a root index and be assigned a priority of “3”. Region **9** may be assigned a default priority. In one embodiment, software may allocate physically contiguous pages for elevated priority regions and pin such pages. Pinning the pages may be performed by, for example, using functions such as “mlock.” Furthermore, the operating system may provide a system call by which application software communicates its dataset priorities to the operating system.

The setting of priorities for memory regions in address spaces **2002**, **2004** may be performed by any suitable mechanism. For example, an instruction for setting cache priority may be defined and available to software. Such an instruction may include a parameter for specifying the process identifier associated with the address space, a start address, an end address, and a priority. For example, region **3** may be set by identifying address space **2002**, a start address of 0xC000, an end address of 0xE000, and a priority of “2”.

In one embodiment, operating system **1808** or another suitable portion of system **1800** may translate the specified virtual memory ranges into physical memory ranges in a lookup table **2008**. In another embodiment, operating system **1808** or another suitable portion of system **1800** may populate a bloom filter table **2006** or other suitable entity. These may be used to populate a suitable entity, such as bloom filter **1842**, for priority address lookup by hardware.

Operating system **1808** may maintain lookup table **2008** has a private data structure for all of the elevated priority physical ranges. Any suitable number of elevated priority ranges may be used. If the priority scheme used by software differs in scale than the priority scheme used by hardware, the priority scheme used by software may be normalized so that the priority values may be correctly read by hardware. Operating system **1808** may maintain  $k$  entries in bloom filter table **2006**, wherein  $k$  is the number of priority levels supported by hardware. Each entry of bloom filter table **2006** may include a bloom filter pattern indexed by the associated priority.

Using the elements of bloom filter table **2006** as programmed into bloom filter **1842**, cache controller **1844** may determine the priority of any given physical address. In one embodiment, priorities may be maintained and looked-up directly in terms of virtual addresses, for example, if caches are accessed by virtual addresses in those embodiments. In another embodiment, cache controller **1844** may perform the look-up in a single clock cycle. The lookup may be specified as yielding a priority  $\pi$  of an address  $\phi$ :

$$\pi(\phi) = \begin{cases} \max(j | \phi \in P_j) \\ 0: \text{ if } \nexists j | \phi \in P_j \end{cases}$$

wherein  $P_i$  is the set of physical addresses mapped to priority  $i$ . Thus, the priority returned is the maximum priority level  $j$  for which the address  $\phi$  may be found within any  $P_j$ . Otherwise, the priority returned is zero. If bloom filter **1842** is used to provide lookup, a statistical error may exist. However, such an error may be kept sufficiently small by choosing a sufficiently large size or width for bloom filter **1842** to avoid error.

Assignment of bloom filter table **2006** to bloom filter **1842** may be performed in any suitable manner. For example, patterns or masks as specified in bloom filter table **1842** may be sent to cache controller **1844**. Furthermore, during runtime the bloom filters may be reassigned as software adjusts priorities. Bloom filters may be loaded at a context switch or may be kept across multiple threads. Such selection may be adjusted dynamically.

FIG. **21** illustrates example operation of system **1800** to perform and adapt cache eviction, according to embodiments of the present disclosure. In one embodiment, a method **2100** may be applied by cache controller **1844** to perform and adapt cache eviction. Thus, method **2100** illustrates example operation of cache controller **1844** or any other suitable part of system **1800**. Method **2100** may begin at any suitable point and may execute in any suitable order. In one embodiment, method **2100** may begin at **2105**.

In one embodiment, cache controller **1844** may utilize priority specified by software in view of a cache victim selection algorithm. Such a cache victim selection algorithm may include any suitable such algorithm, including those discussed above. In another embodiment, if the priority specified by software conflicts with the results of selecting a cache victim (by, for example, selecting a cache victim with a high priority), then cache controller **1844** may repeat the process of selecting a cache victim in an attempt to find a lower-priority cache victim. In yet another embodiment, cache controller **1844** may limit the number of attempts that are performed to find such a lower-priority cache victim and adapt the selection process. In a further embodiment, cache controller **1844** may limit such a number of attempts proportionally to the priority of selected cache victims.

At **2105**, a miss in a cache or other container may be detected. The miss may be for a cacheline or other data with an address  $N$ . At **2110**, cache controller **1844** may determine a candidate cache victim within the cache located at address  $V$  through any suitable cache victim identification algorithm. The candidate cache victim located at address  $V$  may have been, for example, the oldest or last-accessed element within the cache.

At **2115**, cache controller **1844** may determine the priority levels specified for both the requested data at address  $N$  and the candidate cache victim at address  $V$ . In one embodiment, if the candidate cache victim at address  $V$  has an equal or lower priority than the priority of the requested data at address  $N$ , then the candidate may be evicted at **2150**. Otherwise, cache controller **1844** may proceed to **2120** to search for a lower priority candidate to evict.

For example, the requested data at address  $N$  may have a priority of “3” and the eviction candidate at address  $V$  may have a priority of “1.” In such a case, the eviction candidate may be displaced by the requested data in the cache. In another example, the requested data at address  $N$  may have a priority of “1” and the eviction candidate at address  $V$  may



have a priority of “3.” In such a case, cache controller may proceed to determine whether to search for a lower priority candidate to evict.

At **2120**, a maximum number of iterations or another threshold for searching for a lower priority candidate may be established. In one embodiment, the maximum may be set as the previously determined priority level of the requested data at address N. For example, if the requested data at address N has a priority of “3”, then the threshold may be set to three. Moreover, a counter for the number of attempts that cache controller **1844** has made to determine another cache victim of lower priority may be initialized.

At **2125**, cache controller **1844** may identify a new cache victim according to the cache victim identification algorithm. The new cache victim may include an entry in the cache at an address designated as W. The counter for the number of attempts that cache controller **1844** has made to determine another cache victim may be incremented.

At **2130**, cache controller **1844** may access bloom filter **1842** to determine priority for the address designated as W. In one embodiment, cache controller **1844** may also determine priority for the address designated as N and V if such designations have changed. Cache controller **1844** may determine whether the priority of the address designated as W is less than the priority of the previously determined address V. Thus, cache controller **1844** may determine whether the new victim has a lower priority than the previously determined victim. If W does have a lower priority than V, then cache controller **1844** may proceed to **2135** to continue to evaluate W. Otherwise, cache controller **1844** may proceed to **2145** in anticipation of possibly searching for another cache victim. By proceeding with the lowest priority candidate, cache controller **1844** may evict a lowest-priority-available candidate if a fully suitable candidate is not identified.

At **2135**, V may be redesignated and assigned to the newer candidate cache victim at address W. Cache controller **1844** may repeat the above-described evaluation of V at **2140**, wherein cache controller **1844** may determine whether the priority of the candidate cache victim at address V (formerly designated at address W) is less than or equal to the priority of the requested cacheline at address N. If so, cache controller **1844** may proceed to **2150** to evict the data at V for the data at N. Otherwise, cache controller **1844** may proceed to **2145** to determine whether to continue searching for a cacheline with lower priority.

At **2145**, cache controller **1844** may determine whether the count of the number of times that cache controller **1844** has searched for a suitable lower priority cache victim to make room for the identified cacheline at address N has exceeded a threshold. In one embodiment, the threshold may include the threshold specified by the priority of the cacheline at address N, though any suitable threshold may be used. In other embodiments, the threshold may be flexible or adaptable depending upon an average of the number of times lower priority victims are found or not found during previous execution for other cache misses.

In one embodiment, if the threshold has not been exceeded, then cache controller **1844** may find another cache victim to evaluate and return to **2125**. If the threshold has been exceeded, cache controller **1844** may proceed to **2150** and evict the current candidate cache victim (at address V) for the requested cacheline (at address N), even though the current candidate cache victim has a higher priority than the requested cacheline.

Cache controller **1844** may evaluate the number of times it has searched for a replacement cache victim because it is possible that all candidate cache victims might have higher

priority than the requested cacheline. If cache controller **1844** repeats such a search too many times, the performance gains of keeping higher priority data within the cache may be lost.

For example, consider the priority of the requested cacheline N to be “2” and the priority of the first selected cache victim V to be “4.” If at **2130** the priority of the new candidate W is “4” and the priority of the previous candidate V is “3”, then cache controller **1844** may maintain candidate V for subsequent evaluation. If at **2130** the priority of the new candidate W and the priority of the previous candidate V are the same, then cache controller **1844** may maintain candidate V for subsequent evaluation. If at **2130** the priority of the new candidate W is “3” and the priority of the previous candidate V is “4”, then cache controller **1844** may make W the candidate V for subsequent evaluation. In such a case, at **2140**, the priority of V is now “3” which is still greater than the priority of the requested cacheline (“2”), and so cache controller **1844** may continue searching for cache victims, depending upon the analysis in **2145**. If instead the priority of the requested cacheline N was “3”, then the priority of the requested cacheline N would be greater than or equal to the priority of the candidate cache victim V, and thus cache controller **1844** would evict the candidate cache victim V at **2150**.

In another example, if the priority of the requested cacheline N was “2”, for each evaluated cache victim V with a priority greater than “2”, the counter k may be incremented. A maximum value of k may be set to “2” corresponding to the priority of the requested cacheline N. Thus, in this example, cache controller **1844** might only make three attempts to find a cache victim with a lower priority than the requested cacheline N. Accordingly, cache controller **1844** may limit the work performed to find a suitable cache victim in proportion to the priority of the requested cacheline. Furthermore, in this example, if the first three cache victims had a priority greater than “2”, cache controller **1844** would determine at **2145** that k had exceeded the maximum threshold of two attempts to find a suitable cache victim, and would nonetheless replace the presently considered cache victim V with the requested cacheline N. In another example, if the priority of the requested cacheline N were “1”, then cache controller **1844** might only spend a maximum of two iterations searching for a suitable cache victim with a lower priority than the requested cacheline N.

Depending upon the loads being executed and the implementation of system **1800**, cache misses might constitute only a small percent of attempted accesses of memory. Furthermore, a candidate cache victim selected by cache controller **1844** might, more often than not, have a lower priority than the requested cacheline if high-priority designations are used judiciously. Furthermore, the more times that cache controller **1844** looks for additional candidate cache victim, the more likely that a suitable candidate will be found. The likelihood of finding a suitable candidate may be expected to increase exponentially with the number of times that cache controller **1844** searches for a replacement candidate. However, if prioritization as specified by software is incorrect, a poor fit for the loads executing on system **1800**, or if software access patterns have misaligned with the prioritization, then cache controller **1844** may dynamically adjust its search mechanisms for suitable cache victims. Such dynamic adjustment may be made for repeated situations in which cache controller **1844** evicts a higher priority cache victim because cache controller **1844** cannot find a cache victim with a lower priority than the requested cacheline.

FIG. 22 illustrates further example operation of system **1800** to perform and adapt cache eviction, according to embodiments of the present disclosure. System **1800** may

dynamically adjust the search mechanisms used to find suitable cache victims by evaluating searching metrics across multiple searches for replacement cache victims. In one embodiment, a method **2100** may be applied by cache controller **1844** to perform and adapt cache eviction. Thus, method **2200** illustrates example operation of cache controller **1844** or any other suitable part of system **1800**. Method **2200** may begin at any suitable point and may execute in any suitable order. In one embodiment, method **2200** may begin at **2205**.

The operation of cache controller **1844** in FIG. **22** may augment the operation illustrated in FIG. **21**. However, as cache controller **1844** selects a cache victim that is from a high priority dataset, a divergence between specified priority and actual use has occurred. In order to determine whether multiple instances of such divergence are occurring, reflecting a larger problem requiring adaptation, in one embodiment cache controller **1844** may evaluate average of the difference between the number of times a lower priority cacheline has evicted a higher priority entry in the cache and the number of times a lower priority entry in the cache has been evicted for a higher priority cacheline. In another embodiment, cache controller **1844** may also consider a running average of the number of times the cache is searched for a lower priority entry. If cache controller **1844** determines a mismatch in the assigned priority and the execution of cache eviction, cache controller **1844** may adjust its search process and notify software.

At **2205**, a miss in a cache or other container may be detected. The miss may be for a cacheline or other data with an address N. At **2210**, cache controller **1844** may determine a candidate cache victim within the cache located at address V through any suitable cache victim identification algorithm. The candidate cache victim located at address V may have been, for example, the oldest or last-accessed element within the cache. Cache controller **1844** may initialize a counter for the number of attempts, if any, that cache controller **1844** has made to determine another cache victim of lower priority.

At **2215**, cache controller **1844** may determine the priority levels specified for the requested data at address N and the candidate cache victim at address V. In one embodiment, cache controller **1844** may determine whether the candidate cache victim at address V has an equal or lower priority than the priority of the requested data at address N. In another embodiment, the priority of the requested data at address N may be adjusted before it is used. Any suitable mechanism or criteria may be used to adjust the priority of the requested data at address N. In one embodiment, the adjustment may only be made upon a determination that execution in system **1800** has deviated from the patterns indicated in dataset priorities set by software. Otherwise, the unadjusted priority of the requested data at address N may be used. If the candidate cache victim at V has a lower or equal priority as the adjusted priority of the requested cacheline N, then cache controller **1844** may evict the data at V and replace it with the requested cacheline N at **2270**. Otherwise, cache controller **1844** may evaluate possible other candidate cache victims beginning at **2220**.

In one embodiment, the adjustment to priority may be derived from the normal priority and an average number of times that cache controller **1844** has searched for additional candidate cache victims. Such a search may arise when a candidate cache victim has a higher priority than the requested cacheline N, as described above. As this average number rises, the adjusted priority of the requested cacheline may rise in comparison to the real priority. Thus, cache controller makes it more likely for a candidate victim to have a lower priority than the requested cacheline. As a conse-

quence, it is more likely that the candidate cache victim will be evicted and searching for additional candidate cache victims will be lessened.

Any suitable adjustment may be made to priority based upon the average number of times that cache controller **1844** has searched for additional candidate cache victims (referred as “average-k”). In one embodiment, when average-k rises above a threshold, defined in absolute or relative terms, the priority may be increased a defined or percentage amount. As average-k rises above another, higher threshold, the priority may again be increased. In another embodiment, average-k may be normalized and applied as a factor to priority. The resultant priority may be rounded down or up as appropriate. Specific threshold amounts and corresponding priority increases may be determined through experimentation on system **1800**. As average-k decreases below the same thresholds, the priority may be similarly lowered.

For example, if average-k goes above “2”, the priority may be adjusted by adding “1” to the priority of the requested cacheline N. Thus, if the requested cacheline N normally has a priority of “2”, the adjusted priority may be “3”. The adjusted priority is compared to the priority of the candidate cache victim as shown in **2215**. The increased priority of the requested cacheline N makes it more likely that the cacheline N will have greater priority than the candidate cache victim, wherein cache controller **1844** may evict the candidate at **2270**. If average-k goes above “3”, the priority may be adjusted by adding an additional “1” to the priority of the requested cacheline N. Thus, if the requested cacheline N normally has a priority of “2”, the adjusted priority may be “4”. The increased priority of the requested cacheline N makes it even more likely that the cacheline N will have greater priority than the candidate cache victim, wherein cache controller **1844** may evict the candidate at **2270**.

At **2220**, a maximum number of times that cache controller **1844** may search for replacement cache victims may be determined. In one embodiment, such a maximum may be set to the priority of the requested cacheline N, as performed in FIG. **21**. In another embodiment, such a threshold may be adjusted downward when the cache includes so many high priority cache victim candidates that the search for suitable cache victims is impeded. Such a condition may be measured in any suitable manner, such as by average-k. In another embodiment, performing adjustment on such a maximum may be the inverse of the adjustment performed at **2215**. Accordingly, when average-k rises, cache controller **1844** might not spend as many iterations searching for cache victim candidates. The adjustment to the maximum may be reversed, and the maximum raised as average-k reduces to a previous level. Under normal operation, the unadjusted priority of the requested cacheline, and thus an unadjusted maximum number of searches, may be used at **2220**.

Moreover, adjustments to the maximum in **2220** and to the priority of the requested cacheline N in **2215** may be made upon indications received by cache controller **1844** that searches for additional cache victims will not be fruitful. Such situations may include instructions from software that throttle the replacement searches, determinations that a cache is thrashing, determinations that software is going through garbage collection, or any other suitable condition.

At **2225**, cache controller **1844** may identify a new cache victim according to the cache victim identification algorithm. The new cache victim may include an entry in the cache at an address designated as W. The counter for the number of attempts that cache controller **1844** has made to determine another cache victim may be incremented.

At **2230**, cache controller **1844** may access bloom filter **1842** to determine priority for the address designated as W. Cache controller **1844** may determine whether the priority of the address designated as W is less than the priority of the previously determined address V. If W does have a lower priority than V, then cache controller **1844** may proceed to **2135** to continue to evaluate W. Otherwise, cache controller **1844** may proceed to **2245** in anticipation of possibly searching for another cache victim. By proceeding with the lowest priority candidate, cache controller **1844** may evict a lowest-priority-available candidate if a fully suitable candidate is not identified.

At **2235**, V may be redesignated as representing the newer cache victim W. At **2240**, cache controller **1844** may determine whether the priority of the candidate cache victim at address V is less than or equal to the priority of the requested cacheline at address N. In one embodiment, the adjusted priority of N may be used. If so, cache controller **1844** may proceed to **2270** to evict the data at V for the data at N. Otherwise, cache controller **1844** may proceed to **2245** to determine whether to continue searching for a cacheline with lower priority.

In one embodiment, if the new candidate cache victim V has a lower or equal priority than the requested cacheline N, at **2250** cache controller **1844** may determine that a low priority line is to be displaced. This may represent a successful identification and eviction of a lower priority cache victim as foreseen by the designation of priority datasets. Such an action may be tracked as “low-count” across multiple misses and execution of **2200**. Low-count may be incremented at **2250**. If such success become greatly outweighed, on average, by unsuccessful identifications of lower priority cache victims (resulting in eviction of higher priority cache victims), then cache controller **1844** may perform dynamic adjustments or notification of software.

At **2245**, cache controller **1844** may determine whether the count of the number of times that cache controller **1844** has searched for a suitable lower priority cache victim to make room for the identified cacheline at address N has exceeded a threshold. In one embodiment, the threshold may include the threshold specified by the inversely-adjusted priority of the cacheline at address N, though any suitable threshold may be used. If the threshold has been reached, then cache controller **1844** may evict cache victim V at **2270**, even though V has a higher priority than the requested cacheline N. If the threshold has not been reached, cache controller may return to **2225** to select an additional cache victim.

In one embodiment, if the threshold of iterations to find a suitable cache victim has been reached, at **2255** cache controller **1844** may determine that a high priority line is to be displaced. This may represent an unsuccessful identification and eviction of a lower priority cache victim (wherein a higher priority cache victim was evicted instead), in contrast to the designation of priority datasets. Such an action may be tracked as “high-count” across multiple misses and execution of method **2200**. High-count may be incremented at **2255**. If such failures greatly outweigh, on average, the successful identifications of lower priority cache victims, then cache controller **1844** may perform dynamic adjustments or notification of software.

At **2260**, average-k may be updated with the number of times that cache controller searched for additional cache victims during the instant execution of method **2200**. Thus, average-k may reflect a rate of attempted evictions of high priority lines from the cache. At **2265**, it may be determined whether a ratio of or difference between evictions of high cache entries (resulting from **2255**) and evictions of low cache entries (re-

sulting from **2250**) has reached a threshold. The threshold may be, for example, a ratio of 2:1 of evictions of high priority candidates to lower priority candidates. Upon reaching the threshold, cache controller **1844** may make adjustments at **2215** and **2220** to adapt cache eviction, notify software, pause prefetching, or take any other suitable corrective action.

In one embodiment, cache controller **1844** may pause consideration of priority datasets in cache eviction until the ratio or difference between low-count and high-count returns to normal, or at least below the threshold. In such a situation, cache controller **1844** may perform cache victim selection using the algorithm of, for example, **2210**, and then evict the cache victim at **2270** while omitting one or more of the steps of **2215-2245**. For each such eviction, cache controller **1844** may increment low-count, update average-k, and reanalyze the ratio or difference between low-count and high-count.

In another embodiment, cache controller **1844** may identify candidates for cache eviction that included a high priority (designated by the priority datasets) to software. The addresses of such candidates may be logged to, for example, a circular buffer, which can be accessed by operating system **1808** or application **1810** for retrieval. Such software may make adjustments to the priority datasets based upon such identification of badly prioritized memory addresses.

Cache controller **1844** may perform **2250**, **2255**, **2260**, and **2265** in parallel with **2270**. At **2275**, cache controller **1844** may determine whether to continue monitoring for cache misses at **2205** or to terminate.

In one embodiment, cache controller **1844** may reclaim more than one cacheline. Cache controller **1844** may reclaim more than one cacheline with a lower priority than the requested cacheline N.

FIG. **23** is an illustration of operation of software to perform corrective action when notified by hardware that too many high-priority addresses have been evicted from cache, according to embodiments of the present disclosure. Such software performing the operation in FIG. **23** may include, for example, operating system **1808**, application **1810**, a software performance and quality-of-service utility, or DBT **1816**. The operations performed may include a method **2300**. Thus, method **2300** illustrates example operation of software entities or any other suitable part of system **1800**. Method **2300** may begin at any suitable point and may execute in any suitable order. In one embodiment, method **2300** may begin at **2305**.

At **2305**, in one embodiment, software may receive information that priority has been adapted by hardware as described in **2265** of FIG. **23**. In other embodiments, software may determine that action is required based upon executed instructions, a context switch, garbage collection, a phase change, or any other suitable criterion. Software may take any of the corrective action described below alone or together, in series or in parallel.

At **2310**, software may add instructions to inform hardware to pause priority evaluation. In response to such instructions, hardware may perform normal cache victim selection without regard to dataset priorities. The consideration of cache eviction in view of dataset priorities may be resumed upon follow-up instructions by software, a specific time, or other suitable criteria.

At **2315**, software may instruct hardware to resume consideration of dataset priorities after a phase transition, such as a context switch, garbage collection, or other event has completed. In such a case, hardware may have previously stopped considering dataset priorities, whether because of received software instructions or because of its own analysis.

41

At **2320**, software may instruct operating system **1808** to restore values of a previously used bloom filter. Such a bloom filter may have been used in conjunction with a previous thread, and such a restoration may be made upon a context switch wherein the thread will be executed again.

At **2325**, software may adjust dataset priorities. In one embodiment, software may adjust dataset priorities by pruning priority designations corresponding to addresses identified as high-priority candidate cache victims by hardware. Such cache victims may meet the criteria of eviction otherwise applied by hardware, but are high-priority. Software may redesignate such addresses as lower priority or as having no priority. The address ranges may be removed and added again at a later time, according to processor utilization of the addresses.

FIG. **24** is a flowchart of an example embodiment of a method **2400** for executing adaptive dataset priorities, in accordance with embodiments of the present disclosure. Method **2400** may illustrate operations performed by, for example, processor **1804**, applications **1810**, operating system **1808**, DBT **1816**, or cache controllers **1844**. Portions of method **2400** may be performed by portions of methods **2100**, **2200**, or **2300**. Method **2400** may begin at any suitable point and may execute in any suitable order. In one embodiment, method **2400** may begin at **2405**.

At **2405**, priority ranges for cache persistence may be determined for ranges of virtual memory. At **2410**, the priority ranges may be translated to physical memory. Furthermore, a mapping of the priority values may be mapped to the ranges in a bloom filter.

At **2415**, access of an unavailable cache line may be detected. In response to such a cache miss, a potential line to evict may be determined based upon a cache victim algorithm. At **2420**, it may be determined whether the line to be evicted is a high priority line. If not, at **2425** the line may be evicted.

If the line is a high priority line, at **2430** the priority of the cache victim line may be adapted or adjusted, if necessary. At **2435**, it may be determined whether a suitable cache victim has been identified. Such a suitable cache victim may have a priority (possibly adjusted) low enough to be evicted in view of the requested cacheline. Furthermore, if no suitable cache victim has been identified, it may be determined whether the cache victim line will nonetheless be evicted do to operational boundaries. If not, method **2400** may return to **2430**. If so, method **2400** may proceed to **2440**, wherein the cache victim line will be evicted.

At **2445**, if necessary, software may be informed of the operational status of cache evictions. At **2450**, if necessary, software may adjust priority or take other corrective action. At **2455**, it may be determined whether to repeat method **2400**. If so, method **2400** may return to, for example, **2415**. If not, method **2400** may terminate.

Methods **2100**, **2200**, **2300**, and **2400** may be initiated by any suitable criteria. Furthermore, although methods **2100**, **2200**, **2300**, and **2400** describe operation of particular elements, methods **2100**, **2200**, **2300**, and **2400** may be performed by any suitable combination or type of elements. For example, methods **2100**, **2200**, **2300**, and **2400** may be implemented by the elements illustrated in FIGS. **1-20** or any other system operable to implement methods **2100**, **2200**, **2300**, and **2400**. As such, the preferred initialization point for methods **2100**, **2200**, **2300**, and **2400** and the order of the elements comprising methods **2100**, **2200**, **2300**, and **2400** may depend on the implementation chosen. In some embodiments, some elements may be optionally omitted, reorganized, repeated,

42

or combined. Moreover, elements of methods **2100**, **2200**, **2300**, and **2400** may be interchanged or implemented by one another.

Embodiments of the mechanisms disclosed herein may be implemented in hardware, software, firmware, or a combination of such implementation approaches. Embodiments of the disclosure may be implemented as computer programs or program code executing on programmable systems comprising at least one processor, a storage system (including volatile and non-volatile memory and/or storage elements), at least one input device, and at least one output device.

Program code may be applied to input instructions to perform the functions described herein and generate output information. The output information may be applied to one or more output devices, in known fashion. For purposes of this application, a processing system may include any system that has a processor, such as, for example; a digital signal processor (DSP), a microcontroller, an application specific integrated circuit (ASIC), or a microprocessor.

The program code may be implemented in a high level procedural or object oriented programming language to communicate with a processing system. The program code may also be implemented in assembly or machine language, if desired. In fact, the mechanisms described herein are not limited in scope to any particular programming language. In any case, the language may be a compiled or interpreted language.

One or more aspects of at least one embodiment may be implemented by representative instructions stored on a machine-readable medium which represents various logic within the processor, which when read by a machine causes the machine to fabricate logic to perform the techniques described herein. Such representations, known as "IP cores" may be stored on a tangible, machine-readable medium and supplied to various customers or manufacturing facilities to load into the fabrication machines that actually make the logic or processor.

Such machine-readable storage media may include, without limitation, non-transitory, tangible arrangements of articles manufactured or formed by a machine or device, including storage media such as hard disks, any other type of disk including floppy disks, optical disks, compact disk read-only memories (CD-ROMs), compact disk rewritables (CD-RWs), and magneto-optical disks, semiconductor devices such as read-only memories (ROMs), random access memories (RAMs) such as dynamic random access memories (DRAMs), static random access memories (SRAMs), erasable programmable read-only memories (EPROMs), flash memories, electrically erasable programmable read-only memories (EEPROMs), magnetic or optical cards, or any other type of media suitable for storing electronic instructions.

Accordingly, embodiments of the disclosure may also include non-transitory, tangible machine-readable media containing instructions or containing design data, such as Hardware Description Language (HDL), which defines structures, circuits, apparatuses, processors and/or system features described herein. Such embodiments may also be referred to as program products.

In some cases, an instruction converter may be used to convert an instruction from a source instruction set to a target instruction set. For example, the instruction converter may translate (e.g., using static binary translation, dynamic binary translation including dynamic compilation), morph, emulate, or otherwise convert an instruction to one or more other instructions to be processed by the core. The instruction converter may be implemented in software, hardware, firmware,

43

or a combination thereof. The instruction converter may be on processor, off processor, or part-on and part-off processor.

Thus, techniques for performing one or more instructions according to at least one embodiment are disclosed. While certain exemplary embodiments have been described and shown in the accompanying drawings, it is to be understood that such embodiments are merely illustrative of and not restrictive on other embodiments, and that such embodiments not be limited to the specific constructions and arrangements shown and described, since various other modifications may occur to those ordinarily skilled in the art upon studying this disclosure. In an area of technology such as this, where growth is fast and further advancements are not easily foreseen, the disclosed embodiments may be readily modifiable in arrangement and detail as facilitated by enabling technological advancements without departing from the principles of the present disclosure or the scope of the accompanying claims.

What is claimed is:

1. A processor, comprising:

a front end including circuitry to receive an instruction defining a priority dataset, the priority dataset including a plurality of ranges of memory addresses, each range corresponding to a respective priority level;

a cache; and

a cache controller, including circuitry to:

detect a miss in the cache for a requested cache value; determine a candidate cache victim from the cache; determine a priority of the requested cache value according to the priority dataset;

determine a priority of the candidate cache victim according to the priority dataset;

evict the candidate cache victim based on whether the priority of the candidate cache victim is less than or equal to the priority of the requested cache value;

maintain a first count of previous evicted cache victims with higher priority than respective previous requested cache values;

maintain a second count of previous evicted cache victims with lower or equal priority than respective previous requested cache values; and

adjust eviction policies based upon the first count and the second count.

2. The processor of claim 1, wherein the cache controller further includes circuitry to evict the candidate cache victim further based on an instruction to adjust evaluation of the priority dataset.

3. The processor of claim 1, wherein the cache controller further includes circuitry to, upon a determination that the priority of the candidate cache victim is higher than the priority of the requested cache value, search for a new candidate cache victim to compare with the requested cache value.

4. The processor of claim 1, wherein the cache controller further includes circuitry to:

upon a determination that the priority of the candidate cache victim is higher than the priority of the requested cache value, search for a new candidate cache victim with a lower or equal priority to the requested cache value; and

bound the search for the new candidate cache victim with a threshold, the threshold determined from the priority of the requested cache value.

5. A method comprising, within a processor:

receiving an instruction defining a priority dataset, the priority dataset including a plurality of ranges of memory addresses, each range corresponding to a respective priority level;

44

detecting a miss in a cache for a requested cache value;

determining a candidate cache victim from the cache;

determining a priority of the requested cache value according to the priority dataset;

determining a priority of the candidate cache victim according to the priority dataset;

determining whether to evict the candidate cache victim based on whether the priority of the candidate cache victim is less than or equal to the priority of the requested cache value;

maintaining a first count of previous evicted cache victims with higher priority than respective previous requested cache values;

maintaining a second count of previous evicted cache victims with lower or equal priority than respective previous requested cache values; and

adjusting eviction policies based upon the first count and the second count.

6. The method of claim 5, further comprising evicting the candidate cache victim further based on an instruction to adjust evaluation of the priority dataset.

7. The method of claim 5, further comprising, upon a determination that the priority of the candidate cache victim is higher than the priority of the requested cache value, search for a new candidate cache victim to compare with the requested cache value.

8. The method of claim 5, further comprising:

upon a determination that the priority of the candidate cache victim is higher than the priority of the requested cache value, searching for a new candidate cache victim with a lower or equal priority to the requested cache value; and

bonding the search for the new candidate cache victim with a threshold, the Threshold determined from the priority of the requested cache value.

9. A system for executing instructions, including:

a front end including circuitry to receive an instruction defining a priority dataset, the priority dataset including a plurality of ranges of memory addresses, each range corresponding to a respective priority level;

a cache; and

a cache controller, including circuitry to:

detect a miss in the cache for a requested cache value; determine a candidate cache victim from the cache; determine a priority of the requested cache value according to the priority dataset;

determine a priority of the candidate cache victim according to the priority dataset;

evict the candidate cache victim based on whether the priority of the candidate cache victim is less than or equal to the priority of the requested cache value;

maintain a first count of previous evicted cache victims with higher priority than respective previous requested cache values;

maintain a second count of previous evicted cache victims with lower or equal priority than respective previous requested cache values; and

adjust eviction policies based upon the first count and the second count.

10. The system of claim 9, wherein the cache controller further includes circuitry to evict the candidate cache victim further based on an instruction to adjust evaluation of the priority dataset.

11. The system of claim 9, wherein the cache controller further includes circuitry to, upon a determination that the priority of the candidate cache victim is higher than the pri-

45

ority of the requested cache value, search for a new candidate cache victim to compare with the requested cache value.

12. The system of claim 9, wherein the cache controller further includes circuitry to:

upon a determination that the priority of the candidate  
cache victim is higher than the priority of the requested  
cache value, search for a new candidate cache victim  
with a lower or equal priority to the requested cache  
value; and

bound the search for the new candidate cache victim with a  
threshold, the threshold determined from the priority of  
the requested cache value.

\* \* \* \* \*

46